

S U M E X

STANFORD UNIVERSITY
MEDICAL EXPERIMENTAL COMPUTER RESOURCE

RR - 00785

ANNUAL REPORT - YEAR 04

Submitted to
BIOTECHNOLOGY RESOURCES PROGRAM
NATIONAL INSTITUTES OF HEALTH

June, 1977

DEPARTMENT OF GENETICS
STANFORD UNIVERSITY SCHOOL OF MEDICINE
Joshua Lederberg, Principal Investigator

NATIONAL INSTITUTES OF HEALTH
DIVISION OF RESEARCH RESOURCES
BIOTECHNOLOGY RESOURCES PROGRAM

SECTION I - RESOURCE IDENTIFICATION

Report Period:

From August 1, 1976 to July 31, 1977

Grant Number:

RR-00785-04

Report Prepared:

June, 1977

Name of Resource: Stanford University Medical Experimental Computer (SUMEX)	Resource Address: Stanford University Stanford, California 94305	Resource Telephone Number:
Principal Investigator: Joshua Lederberg, Ph.D.	Title: Chairman and Professor Department of Genetics	Academic Department: School of Medicine Department of Genetics
Grantee Institution: Stanford University	Type of Institution: Private University	Investigator's Telephone No.: (415) 497-5801

Name of Institution's Biotechnology Resource Advisory Committee:

SUMEX-AIM Executive Committee

Membership of Biotechnology Resource Advisory Committee:

<u>Name</u>	<u>Title</u>	<u>Department</u>	<u>Institution</u>
Saul Amarel, Ph.D.	Chairman and Professor	Computer Science	Rutgers University
Donald Lindberg, M.D.	Professor Director	Pathology Information Science Group	University of Missouri School of Medicine
Jack Myers, M.D.	University Professor of Medicine	At Large	University of Pittsburgh School of Medicine


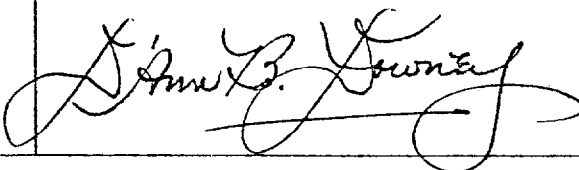
Principal Investigator: Joshua Lederberg, Ph.D. Chairman and Professor	Signature: 	Date: June 10, 1977
Stanford University Official: D'Ann B. Downey Sponsored Projects Officer	Signature: 	Date: June 10, 1977

Table of Contents

BOOK I

Section	Page
Table of Contents - BOOK I	i
List of Figures	ii
Table of Contents - BOOK II	iii
1. RESOURCE OBJECTIVES AND PROGRESS	1
1.1 OVERVIEW OF OBJECTIVES AND RATIONALE	1
1.2 BACKGROUND AND PROGRESS	4
1.2.1 PROGRESS SUMMARY	4
1.2.2 DETAILED PROGRESS REPORT	6
1.2.2.1 DEFINITION OF TERMS AND OBJECTIVES	6
1.2.2.2 FACILITY HARDWARE	7
1.2.2.3 SYSTEM SOFTWARE	14
1.2.2.4 NETWORK COMMUNICATION FACILITIES	16
1.2.2.5 SYSTEM RELIABILITY AND BACKUP	23
1.2.2.6 PROGRAMMING LANGUAGES	23
1.2.2.7 STANFORD AI HANDBOOK PROJECT	26
1.2.2.8 USER SOFTWARE AND INTRA-COMMUNITY COMMUNICATION	27
1.2.2.9 DOCUMENTATION AND EDUCATION	28
1.2.2.10 SOFTWARE COMPATIBILITY AND SHARING	28
1.2.2.11 RESOURCE MANAGEMENT	29
1.2.2.12 SUMMARY OF RESOURCE USAGE	35
1.2.2.13 NETWORK USAGE STATISTICS	48
1.2.2.14 PUBLICATIONS	51

TABLE OF CONTENTS

BOOK I (continued)

List of Figures

1.	SUMEX-AIM Computer Configuration	9
2.	Cost-effectiveness of SUMEX Augmentations	11
3.	Capacity and Loading Increase with Dual Processor Augmentation	13
4.	TYMNET Network Map	19
5.	ARPANET Geographical Network Map	20
6.	ARPANET Logical Network Map	21
7.	Monthly CPU Time Consumed	35
8.	CPU Usage by Community	37
9.	File Space Usage by Community	38
10.	Average Diurnal Loading (3/77): Total Number of Jobs	45
11.	Average Diurnal Loading (3/77): Percent Time Used	45
12.	Average Diurnal Loading (3/77): Percent Overhead	46
13.	Average Diurnal Loading (3/77): Balance Set - Jobs in Core	46
14.	Average Diurnal Loading (3/77): Runnable Jobs	47
15.	TYMNET and ARPANET Usage Data	49

Table of Contents

BOOK II

Introduction	1
5. COLLABORATIVE PROJECT PROGRESS AND OBJECTIVES	41
6.1 STANFORD PROJECTS	41
6.1.1 DENDRAL PROJECT	42
6.1.2 HYDROID PROJECT	76
6.1.3 MOLGEN PROJECT	81
6.1.4 MYCIN PROJECT	84
6.1.5 PROTEIN STRUCTURE PROJECT	108
6.2 NATIONAL AIM PROJECTS	112
6.2.1 ACQUISITION OF COGNITIVE PROCEDURES (ACT)	113
6.2.2 CHEMICAL SYNTHESIS PROJECT (SECS)	118
6.2.3 HIGHER MENTAL FUNCTIONS PROJECT	128
6.2.4 INTERNIST PROJECT	132
6.2.5 MEDICAL INFORMATION SYSTEMS LABORATORY	138
6.2.6 RUTGERS COMPUTERS IN BIOMEDICINE	144
6.3 PILOT STANFORD PROJECTS	158
6.3.1 GENETICS APPLICATIONS PROJECT	159
6.3.2 BAYLOR-METHODIST CEREBROVASCULAR PROJECT	161
6.3.3 COMPUTER ANALYSIS OF CORONARY ARTERIOGRAMS	165
6.3.4 QUANTUM CHEMICAL INVESTIGATIONS	169
6.4 PILOT AIM PROJECTS	171
6.4.1 COMMUNICATION ENHANCEMENT PROJECT	172
6.4.2 AI IN PSYCHOPHARMACOLOGY	179
6.4.3 ORGAN CULTURE PROJECT	189
6.4.4 NEUROPROSTHESES PROJECT	191
6.4.5 MATHEMATICAL MODELING OF PHYSIOLOGICAL SYSTEMS	194
6.4.6 PUFF/VM PROJECT (Abridged - Proposal pending)	197

TABLE OF CONTENTS

BOOK II (continued)

Appendix I

OVERVIEW OF ARTIFICIAL INTELLIGENCE RESEARCH	202
[Deleted - Copies available on request]	

Appendix II

AI HANDBOOK OUTLINE	225
-------------------------------	-----

Appendix III

SUMMARY OF MAINSAIL LANGUAGE FEATURES	231
---	-----

Appendix IV

MICROPROGRAMMED MAINSAIL PLANS	235
--	-----

Appendix V

AIM MANAGEMENT COMMITTEE MEMBERSHIP	239
---	-----

Appendix VI

USER INFORMATION - GENERAL BROCHURE	243
---	-----

Appendix VII

GUIDELINES FOR PROSPECTIVE USERS	245
--	-----

BOOK I

SUMEX-AIM Resource Progress Report

This is an annual report of the work performed under an NIH Biotechnology Resources Program grant supporting the Stanford University Medical EXperimental computer (SUMEX) research resource for applications of Artificial Intelligence in Medicine (AIM). It spans the year from May 1976 - April 1977. As we have invested substantial effort in preparing a related document, an application for renewal dated June 1, 1977, this report has been prepared by revising and augmenting the other. Some sections may inadvertently reflect that provenience, e.g., by adopting a longer time perspective, but we believe without distorting or misrepresenting our last year's effort. Book II of this report is the same text as used for the renewal, and contains detailed progress reports of collaborating user projects and other pertinent appendices.

1 RESOURCE OBJECTIVES AND PROGRESS1.1 OVERVIEW OF OBJECTIVES AND RATIONALE

The SUMEX-AIM project is a national computer resource with a dual mission: 1) the promotion of applications of artificial intelligence (AI) computer science research to biological and medical problems and 2) the demonstration of computer resource sharing within a national community of health research projects.

Definitive funding of the SUMEX-AIM resource was initiated in December 1973. The principal hardware was delivered and accepted in April 1974, and the system became operational for users during the summer of 1974. The present renewal is therefore written from a perspective of just short of three years of experience in attempting to develop and serve the user community for the resource.

The original SUMEX proposal was an outgrowth of two lines of endeavor at Stanford that had been supported by the Biotechnology Resources Program. The ACME project (Advanced Computer for MEDical Research), 1965-72, had introduced the innovation of interactive time-shared computing to the medical research community at the Stanford Medical Center. The second line, the DENDRAL project, is a resource-related project connected with applications of artificial intelligence to problems of molecular characterization by analytical instruments like mass-spectrometry, gas-chromatography, nuclear magnetic resonance, and so on.

In 1972 we applied to NIH for the establishment at Stanford of a next generation computer resource to supplant ACME for applications for which the university-wide facility was inadequate. The DENDRAL project was the central source of this initiative; several others entailing real-time instrumentation as much as AI needs were also specified. During the subsequent 18 months, we entered a phase of protracted review and negotiations with BRP and its advisory

groups, from which emerged the policy determination that resources of this scope were best justified if they could be functionally specialized, but geographically generalized. The emerging technology of computer networking opened an opportunity to demonstrate this model in a way that could serve both local and national needs.

Our technical task has been achieved: to collect and implement an effective set of hardware and software tools supporting the development of large and complex AI programs and to facilitate communications and interactions between user groups. In effect, users throughout the country can turn on their own teletype or CRT-display terminals, dial a local number, and logon to SUMEX-AIM with the same ease as if it were located on their own campus -- and have access to a specialized resource unlikely to be matched nearby. From the community viewpoint, we have substantially increased the roster of user projects (from an initial 5) to 11 current major projects plus a group of pilot efforts. Many of these projects are built around the communications network facilities we have assembled; bringing together medical and computer science collaborators from remote institutions and making their research programs available to still other remote users. As discussed in the sections describing the individual projects, a number of the computer programs under development by these groups are maturing into tools increasingly useful to the respective research communities. The demand for production-level use of these programs has surpassed the capacity of the present SUMEX facility and has raised the general issues of how such software systems can be optimized for production environments, exported, and maintained.

1.2 BACKGROUND AND PROGRESS

1.2.1 PROGRESS SUMMARY

This progress summary covers the period from December 1973, when the SUMEX-AIM resource was initially funded, through April 1977. During this period we have met all of the defined goals of the resource:

- i) We have established an effective computing facility to support a nation-wide community of medical AI research projects including connections to two computer communication networks to provide wide geographical access to the facility and research programs.
- ii) We have actively recruited a growing community of user projects and collaborations. The initial complement of collaborators included five projects. This roster has grown to eleven fully authorized projects currently plus a group of approximately six pilot efforts in various stages of formulation. Recruiting efforts have included a public dedication and announcement of the resource, NIH referrals from computer-based project reviews, direct contacts by resource personnel and on-going projects as well as contacts through the AIM workshop series coordinated by the Rutgers Computers in Biomedicine resource under Dr. Saul Amarel.
- iii) We have established an AIM community management structure based on an overseeing Executive Committee and an Advisory Group to assist in recruiting and assessing new project applications and in guiding the priorities for SUMEX-AIM developments and resource allocations. These committees also provide a formal mechanism for user projects to request adjustments in their allocated share of facility resources and to make known their desires for resource developments and priorities.
- iv) SUMEX user projects have made good progress in developing more effective consultative computer programs for medical research; one of the major goals toward which our AI applications are aimed. These performance programs provide expertise in analytical biochemical analyses and syntheses, medical diagnoses, and various kinds of cognitive and affective psychological modeling.
- v) We have worked hard to build system facilities to enable the inter- and intra- group communications and collaborations upon which SUMEX is based. We have a number of examples in which user projects combine medical and computer science expertise from geographically remote institutions and numerous examples of users from all over the United States and occasionally from Europe experimenting with the developing AI programs. The SUMEX staff itself has had good success in establishing such sharing relationships on a system level with other research groups and has many examples of complementary development and maintenance agreements for system programs.
- vi) We have made numerous improvements to the computing resource to extend its capacity, to improve its efficiency, to enhance its human interfaces, to improve its documentation, and to enhance the range of software facilities available to user projects.

- vii) We have begun a core research effort to investigate alternatives and programming tools to facilitate the exportability of user and system software. This is just now producing a "machine-independent" implementation of the ALGOL-like SAIL language which will run on a range of large and small machines and provide a language base for transferring programs.
- viii) We have supported community efforts in the more systematic documentation of AI concepts and techniques and in building more general software tools for the design and implementation of AI application programs. These have included a Stanford AI Handbook project comprising a compendium of short articles about the projects, ideas, problems, and techniques that make up the field of AI.

1.2.2 DETAILED PROGRESS REPORT

The following material covers in greater detail the SUMEX-AIM resource activities over the past 3.5 years. These sections attempt to define in more detail the technical objectives of our research community and include progress in the context of the resource staff and the resource management. Details of the progress and plans for our external collaborator projects are presented in Section 6 on page 41 (in Book II).

1.2.2.1 DEFINITION OF TERMS AND OBJECTIVES

Artificial Intelligence is a branch of computer science which attempts to discern the underlying principles involved in the acquisition and utilization of knowledge in reasoning, deduction, and problem-solving activities (1). Currently authorized projects in the SUMEX community are concerned in some way with the application of these principles to biomedical research. The tangible objective of this approach is the development of computer programs which, using formal and informal knowledge bases together with mechanized hypothesis formation and problem solving procedures, will be more general and effective consultative tools for the clinician and medical scientist. The exhaustive search potential of computerized hypothesis formation and knowledge base utilization, constrained where appropriate by heuristic rules or interactions with the user, has already produced promising results in areas such as chemical structure elucidation and synthesis, diagnostic consultation, and mental function modeling. Needless to say, much is yet to be learned in the process of fashioning a coherent scientific discipline out of the assemblage of personal intuitions, mathematical procedures, and emerging theoretical structure of the "analysis of analysis" and of problem solving. State-of-the-art programs are far more narrowly specialized and inflexible than the corresponding aspects of human intelligence they emulate; however, in special domains they may be of comparable or greater power, e.g., in the solution of formal problems in organic chemistry or in the integral calculus.

An equally important function of the SUMEX-AIM resource is an exploration of the use of computer communications as a means for interactions and sharing between geographically remote research groups in the context of medical computer science research. This facet of scientific interaction is becoming increasingly important with the explosion of complex information sources and the regional specialization of groups and facilities that might be shared by remote researchers. Our community building role is based upon the current state of computer communications technology. While far from perfected, these new capabilities offer highly desirable latitude for collaborative linkages, both within a given research project and among them. Several of the active projects on SUMEX are based upon the collaboration of computer and medical scientists at

(1) For recent reviews to give some perspective on the current state of AI, see: (i) Winston, P.H., "Artificial Intelligence", Addison-Wesley Publishing Co., 1977; (ii) Nilsson, N.J., "Artificial Intelligence", Information Processing 74, North-Holland Pub. Co. (1975); and (iii) a summary by Feigenbaum, E. A., attached as Appendix I, page 202 (see Book II). An additional overview of research areas in AI is provided by the outline for an "Artificial Intelligence Handbook" being prepared under Professor Feigenbaum by computer science students at Stanford (see Appendix II on page 225 in Book II).

geographically separate institutions; separate both from each other and from the computer resource. The network experiment also enables diverse projects to interact more directly and to facilitate selective demonstrations of available programs to physicians and medical students. Even in their current developing state, we have been able to demonstrate that such communication facilities allow access to the rather specialized SUMEX computing environment and programs from a great many areas of the United States (even to a limited extent from Europe) for potential new research projects and for research product dissemination and demonstration. In a similar way, the network connections have made possible close collaborations in the development and maintenance of system software with other facilities.

1.2.2.2 FACILITY HARDWARE

Based on the AI mission of SUMEX-AIM, we selected a Digital Equipment Corporation (DEC) model KI-10 computer system for our facility. This selection was based on 1) hardware architectural and performance features, 2) available software support relevant to AI applications, 3) price versus performance data for the system, and 4) the scope of the user community from which we might expect to draw collaborators and share software. This choice has proved highly effective.

The current system hardware configuration is diagrammed in Figure 1 on page 10. It is the result of a number of augmentations over the past 3 years to meet the capacity needs of the growing SUMEX-AIM project community. Our initial configuration consisted of a KI-10 processor, core memory (192K 36-bit words @ 1 microsecond), swapping storage (1.7M words @ 8 msec average rotational latency and 2 microsecond/word transfer rate), file storage (40M words), magnetic tapes, DEC tapes, terminal line scanner, and line printer. Our network connections are discussed in Section 1.2.2.4 on page 16.

This system reached prime-time saturation by fall of 1974. Since many of our medical and other professional collaborators cannot adjust their schedules to match light computer loading during the night-time hours, the prime-time responsiveness is crucial to being able to support medical experimentation with developing programs and to allow community growth. We have taken active steps to transfer as much prime-time loading as feasible to evening and night hours including shifting personnel schedules (particularly for Stanford-based projects), controlling the allocations of CPU resources between various user communities and projects, and encouraging jobs not requiring intimate user interaction to run during off hours by developing batch job facilities. Despite these efforts, prime-time loading has remained quite high, particularly with the growth of the number of user projects.

A similar congestion has persisted in the on-line file space we have been able to allocate to user projects. Again we have implemented controls to try to assure effective use of available space and to encourage use of external file storage facilities such as the ARPANET Data Computer and other computer sites. Nevertheless, the interactive character of SUMEX use, the large AI program files, and the extensive use of SUMEX for collaborator communications have continuously raised file space demands beyond those we could meet.

We have proposed a number of hardware configuration augmentation steps to the Executive Committee to cost-effectively provide additional capacity. These were based on analyses of predominant system bottlenecks and enhancement steps feasible within available budgets. The enhancements approved by the committee and implemented include:

- 1) Add 64K words of core memory and 20M words of file storage (11/74)
- 2) Add second KI-10 CPU for dual processor operation (5/76)
- 3) Add 256K words of core memory and upgrade file system to higher volume, lower cost technology (recently approved by NIH and the AIM Executive Committee with implementation in progress)

A plot of effective CPU capacity as a function of continuing investment is shown in Figure 2 on page 11 and displays the cost-effectiveness of our sequential augmentations. At the present time our hardware configuration has grown about as much as is cost-effective. Additional growth would entail significant redesigns of the system including upgrades of existing hardware. Contemplating such future expansion also raises the issues of compatibility with newer hardware technologies being announced. These provide advantages in speed, cost, size, and maintainability. Such a complete upgrade is not envisioned in the immediate future as a number of interesting new product announcements are expected over the next 1 or 2 years that could substantially affect such an upgrade strategy.

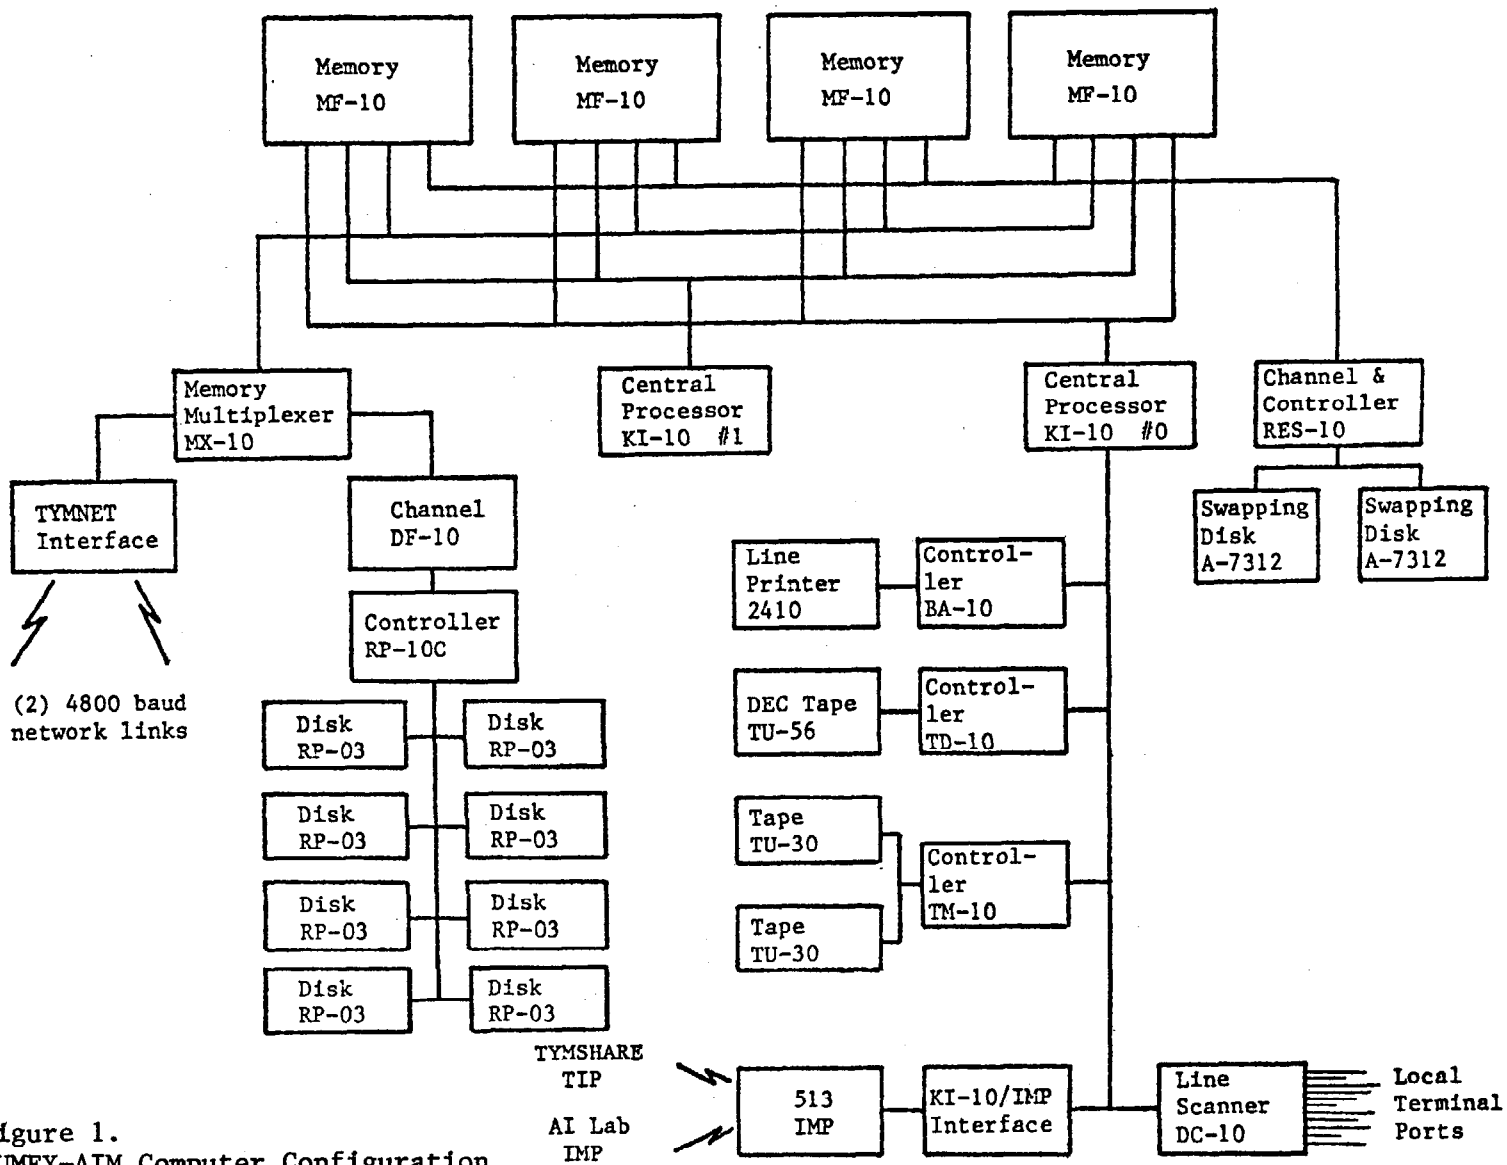



Figure 1.
SUMEX-AIM Computer Configuration

TYMSHARE
TIP
AI Lab
IMP

Figure 2. Cost-effectiveness of SUMEX Augmentations

Estimated Capacity in
Useful KI-10 Equivalents
(Net of overhead)



This plot illustrates the incremental increases in computing capacity achieved as a function of cumulative investment in the SUMEX-AIM facility. The higher slope of the curve after the initial investment illustrates both the substantial investment in peripheral devices (file system, tapes, communications, etc.) and the trend toward lower memory prices. The largest impact in terms of PDP-10 memory price reductions occurred around the time of adding the 64K increment in November 1974. Since then processor prices have stayed relatively stable and memory prices have dropped less dramatically. It should be noted that semi-conductor memories have not yet made a big in-road in the PDP-10 market; this technology is where the more recent memory price reductions have occurred.

The original purchase of 1 KI-10 with 192K of memory for about \$800K performed with about 60% efficiency under peak load. Adding the 64K of memory for \$75K brought the efficiency up to about 85%. Then adding the second processor for \$200K increased throughput to about 1.3-1.4 KI-10 equivalents. This step represents about a 59% increase in throughput for a 20% increased investment. A proposal has been approved recently by the AIM Executive Committee and NIH to augment core memory by 256K words. This augmentation would increase throughput to about 1.7 KI-10 equivalents for another \$100K; this would be a 26%

throughput increase for 8% additional investment. As part of the proposed memory augmentation we plan to upgrade the file and tape systems as well to relieve file space congestion and increase system operations efficiency. Including the net cost of the file/tape upgrade in these figures (purchase price less resale of existing equipment) raises the proposed additional investment to \$160K and the fractional increase from 8% to 13%. Of course, the disk upgrade affects CPU throughput only indirectly in that the increased speed reduces contention, particularly when moving head swapping is necessary. It contributes primarily to supporting the growing on-line file needs of the projects.

Figure 3. Capacity and Loading Increase with Dual Processor Augmentation

	1-PROC OP'N 1/76 - 4/76 -----	2-PROC TRNS'N 5/76 - 8/76 -----	2-PROC OP'N 9/76 - 12/76 -----	2-PROC OP'N 1/77 - 3/77 -----
Peak Ld Ave	4.8	5.6	6.0	6.6
Peak Jobs	30.2	33.3	34.7	38.1
% Overhead/ Processor	18.1	31.1	33.2	31.9
Total CPU Hrs/Mo	304.4	384.9	534.0	520.1

This table presents system usage data averaged over several months preceding, during, and after installation of the SUMEX-AIM dual processor system in order to show real changes in peak loading capacity and computing resources delivered. The first three rows of data are derived from monthly diurnal loading data and reflect average prime-time peak loading conditions (daily peak usage figures are often considerably higher, but those shown better represent gross trends). The last row gives average total monthly CPU hours delivered during the various periods.

With the common criterion that users have pushed both the single and dual processor systems to the limits of useful work in terms of prime time responsiveness, it is clear that the second processor has substantially increased throughput ("tolerable" peak load average up 38%, number of jobs up 26%, and delivered CPU hours up 71%). At the same time the overhead burden per machine has risen from 18 to 32%, principally in the category of I/O wait (total scheduler time and time waiting for a runnable job to be loaded in core). An additional factor, not explicitly shown in these data (because we only have a 1 msec clock), is the added time spent at interrupt level servicing drum swapping. This adds another 10-15% estimated overhead.

We feel these increased overhead figures can be reduced roughly to the single processor levels by adding more memory, thereby effectively recovering about 40-50% of the capacity of a KI-10 processor. A proposal is now pending with the AIM Executive Committee for this augmentation and we expect it to be implemented within the funding ceiling of the current grant.

1.2.2.3 SYSTEM SOFTWARE

In parallel with the choice of DEC PDP-10 hardware for the SUMEX-AIM facility, we selected the TENEX operating system developed by Bolt, Baranek, and Newman (BBN) as the most effective for our medical AI applications work. TENEX was the only available demand-paged system to support simultaneous large address space users, offered the INTERLISP language for LISP-oriented program development, and was well integrated with the ARPANET facilities which provide an excellent base for our community sharing efforts. This choice has proven a very effective one in that the productivity of the TENEX community in AI research has been highly advantageous to us (2).

The original BBN TENEX was written for a hardware-modified KA-10 system. This version of the system required a substantial amount of work to accommodate the relatively limited paging facilities of the KI-10 to run effectively. These early phases also included substantial monitor work to incorporate the TYMNET memory-sharing interface which connects us to the TYMNET and to integrate the high speed swapping storage. We have made numerous enhancements to the monitor calls and corrections of bugs to develop a highly reliable and effective operating system for our community work.

We continue to work to improve the efficiency of the system and its effectiveness in allocating valuable resources. For example we have modified the handling of user page tables so that the expensive procedure of clearing page tables and setting them up to run time-shared users could be minimized. This involved creating a pool of page tables which could be allocated to currently running users and could be kept available without setup overhead. We also implemented a system for migrating dormant pages from our fast swapping storage to moving head disk. This preserves the use of this limited resource for the currently active jobs.

We have implemented a form of "soft" CPU allocation control in the monitor, assisted by a program which adjusts user percentages for the scheduler based on the dynamic loading of the system. The allocation control structure works based on the scheduler queue system and takes account of the a priori allocation of CPU time and that actually consumed. Our TENEX uses a hierarchy of five queues for jobs ranging from highly interactive jobs requiring only small amounts of CPU time between waits to more CPU intensive jobs which can run for long periods without user interaction. These interactive queues (text editing, etc.) are scheduled at highest priority without consideration of allocation percentages. If nothing is runnable from the high priority queues, the CPU-bound queues are scanned and jobs are selected for running based on how much of their allocated time has been received during a given allocation cycle time (currently 100 seconds). If no such jobs are runnable, then those that have received their allocation of CPU time already are scheduled based on how much they are over allocation and how long they have waited to be run again. This system is not a reservation system in that it does not guarantee a given user some percentage of

(2) It should be noted that DEC has recently adopted a form of TENEX (TOPS-20) as their choice for future system marketing. They have made improvements in a number of areas of the monitor and subsystem software but have also shown an increasing tendency to make changes to the TOPS-20 system that impair compatibility with older TENEX systems.

the system. It allocates cycles preferentially, trading off a priori allocations with actual demand but does not waste cycles. This allocation control system is still in an experimental state and we are attempting to evolve the "best" policies with the AIM Executive Committee for dividing the system fairly and effectively among the various communities of users.

During the spring of 1976 we implemented a dual processor version of TENEX as the most cost-effective way to increase our processing capacity. In order to upgrade to the new KL-"n" technology, we would have had to replace most of the equipment that had been purchased initially. For the cost of an additional processor and 8 man-months of intensive software development we were able to increase our CPU capacity by 75%. We have an additional 40% equivalent of a KI-10 processor which can be made available by increasing memory to reduce our swapping contention. The dual processor system that has evolved is running quite reliably. It treats the two machines in an almost symmetric manner. The only difference is that one of the machines has all of the I/O equipment attached to it. They both schedule jobs independently and share the rest of the non-I/O-device monitor code. The areas of the monitor involving the management of resources and jobs which cannot be manipulated by both machines simultaneously are protected by a system of locks. We have made some measurements indicating that overhead for lock waits is less than 10%. The overall increase in capacity provided by the processor upgrade is illustrated in Figure 3 on page 13 which measures key loading parameters in the periods before and after the dual processor installation. Observing the delivery of DEC's high-performance KL-TENEX systems over the past 6 months, it seems clear that for the investment, we made the best choice for the community by implementing the dual processor upgrade. We hope to augment the memory soon to finish exploiting the capacity this extra machine provides and to remove some non-linearities remaining in system swapping performance.

Now that the dual processor system has stabilized, we are undertaking another assessment of system performance to be sure we have removed residual and correctable inefficiencies. This study is on-going now.

Finally, over the past year we made several substantial improvements in the "GTJFN" monitor call which interactively acquires handles on file names specified by the user. These extensions allow for more general "wild card" specifications and interactive help in deciding between and searching for existing file name alternatives. They also give the user much more flexibility in designating groups of files and therefore in structuring his data.

With a working dual processor system, the current implementation of allocation controls in our system, the diverging path of the DEC TOPS-20 system, the termination of active BBN TENEX development, and the unique complications of the KI-10 paging system, we have not made any concerted effort to upgrade our TENEX system to the latest BBN release (1.34). The advantages of such an upgrade are not overwhelming in face of the complicated conversion (KI paging, dual processor, special swapping device handler, TYMNET service routines, local JSYS's, etc.) and resulting system unreliability for some period.

Another area of software development is in the EXECutive program which is the basic user interface to manipulate files, directories, and devices; control job and terminal parameter settings; observe job and system status; and execute public and private programs. This work improves system accommodation to users and provides more convenient and useful information about system and job status. Through such features as login default files, directed file search path commands, mail notification, help facilities, better file archival and retrieval commands, and flexible status information, we have tried to make it easier for users to work on the SUMEX-AIM machine.

1.2.2.4 NETWORK COMMUNICATION FACILITIES

A highly important aspect of the SUMEX system is effective communication with remote users. In addition to the economic arguments for terminal access, networking offers other advantages for shared computing such as uniform user access to multiple machines and special purpose resources, convenient file transfers for software sharing and multiple machine use, more effective backup, co-processing between remote machines, and improved inter-user communications. Over the past year we have been substantially aided in exporting the MAINSAIL system through our network connections. Because of the developmental nature of the language at present, it is important that we have close interactions with the user community and that we be able to effectively perform bug fixes and upgrades. Since MAINSAIL by its nature involves operations on a variety of machines and since our access to example systems cannot be entirely local, the network connections to Rutgers, the Stanford AI Lab, and Stanford Research Institute have been invaluable. It would be considerably more difficult to export MAINSAIL and communicate with users via tapes and mail.

We have based our remote communication services on two networks - TYMNET and ARPANET. These were the only networks existing at the start of the project which allowed foreign host access. Since then, other commercial network systems (notably TELENET) have come into existence and are growing in coverage and services. The two networks to which we are currently connected complement each other; the TYMNET providing primarily terminal service with very broad geographical coverage and unrestricted user access, and the ARPANET having more limited access but providing a broader range of communication services. Together, these networks give a good view of the current strengths and weaknesses of this approach.

Users asked to accept a remote computer as if it were next door will use a local telephone call to the computer as a standard of comparison. Current network terminal facilities do not quite accomplish the illusion of a local call. Data loss is not a problem in network communications - in fact with the more extensive error checking schemes, data integrity is much higher than for a long distance phone link. On the other hand, networking relies upon shared community use of telephone lines to procure widespread geographical coverage at substantially reduced cost. However, unless enough total line capacity is provided to meet peak loads, substantial queueing and traffic jams result in the loss of terminal responsiveness.

TYMNET:

Networks such as TYMNET are a complex interconnection of nodes and lines spanning the country (see Figure 4 on page 20). The primary cause of delay in passing a message through the network is the time to transfer a message from node to node and the scheduling of this traffic over multiplexed lines. This latter effect only becomes important in heavily loaded situations; the former is always present. Clearly from the user viewpoint, the best situation is to have as few nodes as possible between him and the host - this means many interconnecting lines through the network and correspondingly higher costs for the network manager. TENEX in some ways emphasizes this conflict more than other time-sharing systems because of the highly interactive nature of terminal handling (e.g., command and file name recognition and non-printing program commands as in text editors or INTERLISP). In such instances, individual characters must be seen by the host machine to determine the proper echo response in contrast to other systems where only "line at a time" commands are allowed. We have connected SUMEX to the TYMNET in two places as shown in Figure 4 so as to allow more direct access from different parts of the country. Based on delay time statistics collected during the previous year from our TYMSTAT program, the response times are scarcely acceptable. When delay times exceed 200-300 milliseconds, the character printing lag problems become noticable with a full duplex, 30 char/sec terminal. In the past these times have been particularly bad in New York with peak delays approaching 3 seconds one way! Other nodes have shown uniformly high readings as well. These data were reflected in the subjective, but strongly articulated, comments of many of our user groups.

We have had numerous meetings with TYMNET personnel to try to ease these problems and have instituted reroutings of the lines connecting SUMEX-AIM to the network. Also local lines to more strategic terminal nodes have been considered for users in areas poorly served by the existing line layout. TYMNET has also made some upgrades in the internal connectivity and speeds with which data is switched within their node clusters. These changes seem to have had some beneficial effects in that delay times have improved and user complaints have subsided.

We will continue to pursue improvements in TYMNET response but user terminal interactions such as used in TENEX programs are not realized in the time-sharing systems offered by most other TYMNET users and hence are not supported well by TYMNET. TYMNET has implemented 1200 baud service in 7 major cities over the past year. Unfortunately many of our users are not in these cities so we have only limited experience with the 1200 baud support.

ARPANET:

The ARPANET, while designed for more general information transfer than purely terminal handling, has similar bottleneck problems in its topology (see the current geographical and logical maps of the ARPANET in Figure 5 and Figure 6 on page 21). These are reduced by the use of relatively higher speed interconnection lines (50 K baud instead of 2400 - 9600 baud lines as in TYMNET) but response delays through many nodes become objectionable eventually as well.

Consistent with the agreements with ARPA when we were granted network access initially, we are enforcing a policy to restrict the use of the ARPANET to users who have affiliations with ARPA-supported contractors and system/software interchange with cooperating TENEX sites. The administration of the network passed from the ARPA Information Processing Techniques Office to the Defense Communications Agency as of July 1975. At that time policies were announced restricting access to DoD-affiliated users. We have restricted the facilities for calling from SUMEX out to other sites on the ARPANET to authorized users. This also protects the SUMEX-AIM machine from acting as an expensive terminal handler for other machines - this function is better fulfilled by dedicated terminal handling machines (TIPS). In general, we have developed excellent working relationships with other sites on the ARPANET for system backup and software interchange - such day-to-day working interactions with remote facilities would not be possible without the integrated file transfer, communication, and terminal handling capabilities unique to the ARPANET.

We take very seriously the responsibility to provide effective communication capabilities to SUMEX-AIM users and are continuously looking for ways to improve our existing facilities as well as investigate alternatives becoming available. We have done preliminary investigations of the TELENET facilities that have been rapidly expanding this past year. BB&N has hooked one of their TENEX systems up to TELENET and whereas we did not have the same quantitative tools we have for measuring response on the TYMNET, we observed TELENET delays at least as long as those encountered on TYMNET. We did the reverse experiment by using long distance telephone to connect from the TELENET node in Washington, D.C. to the SUMEX machine in California and observed the same sort of delays reaching several seconds per character. The TELENET has many attractive feature in terms of a symmetry analogous to that of the ARPANET for terminal traffic and file transfers and being commercial would not have the access restrictions of the ARPANET. However, until the network throughput improves we would not get substantial benefits from connecting to it.

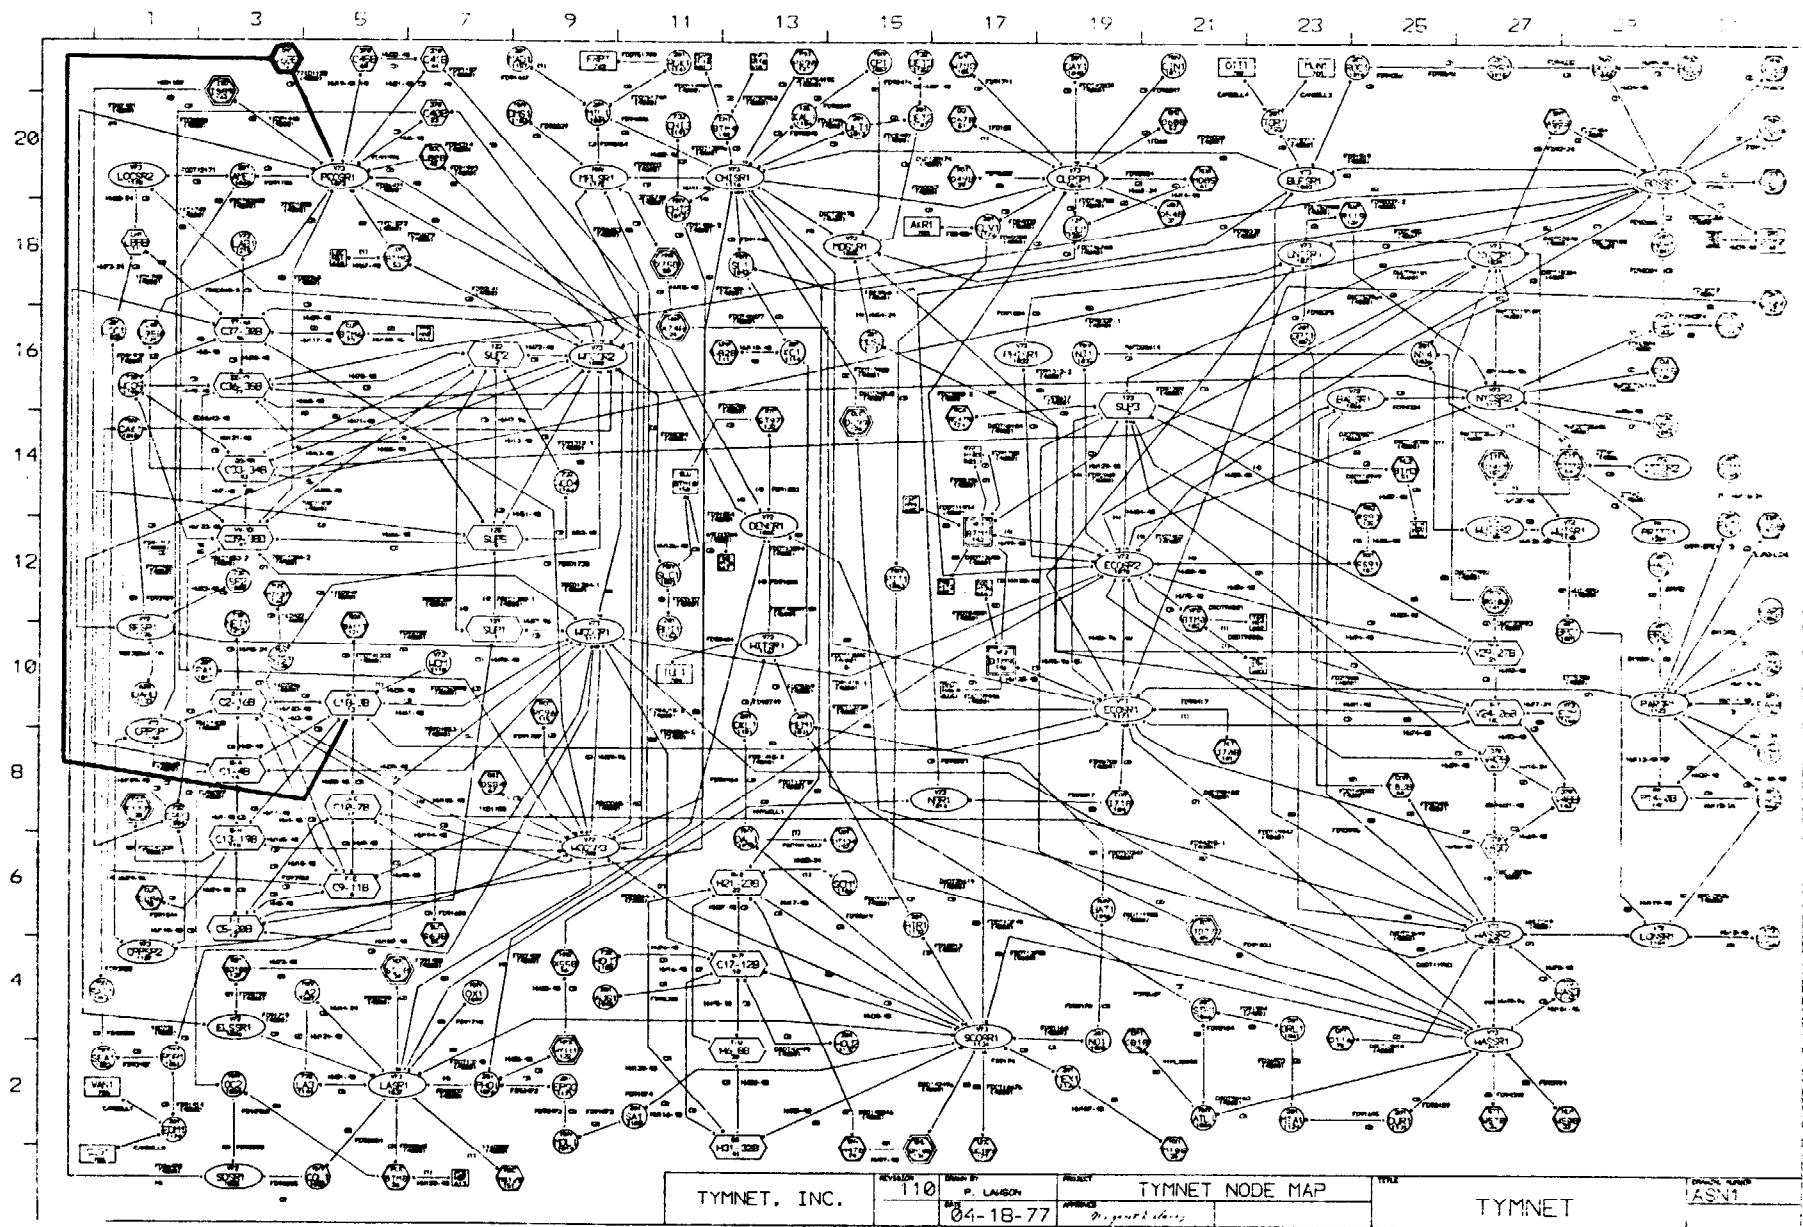


Figure 4. TYMNET Network Map

Figure 5.

ARPANET GEOGRAPHIC MAP, APRIL 1977

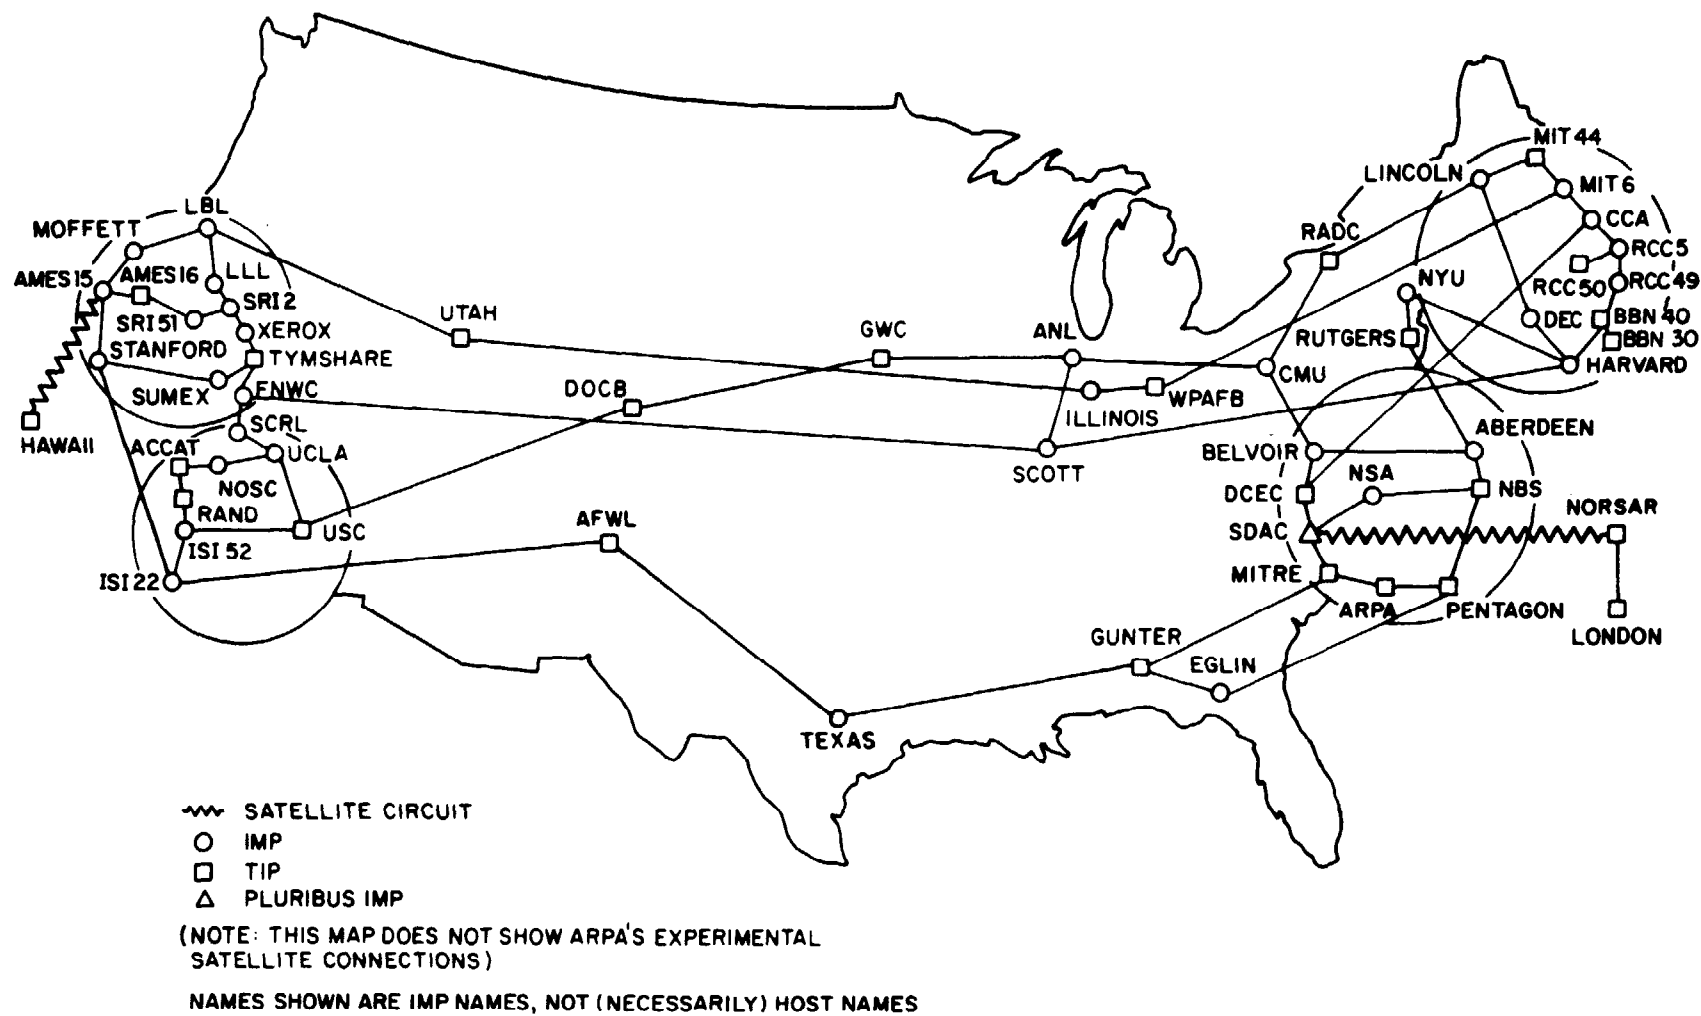
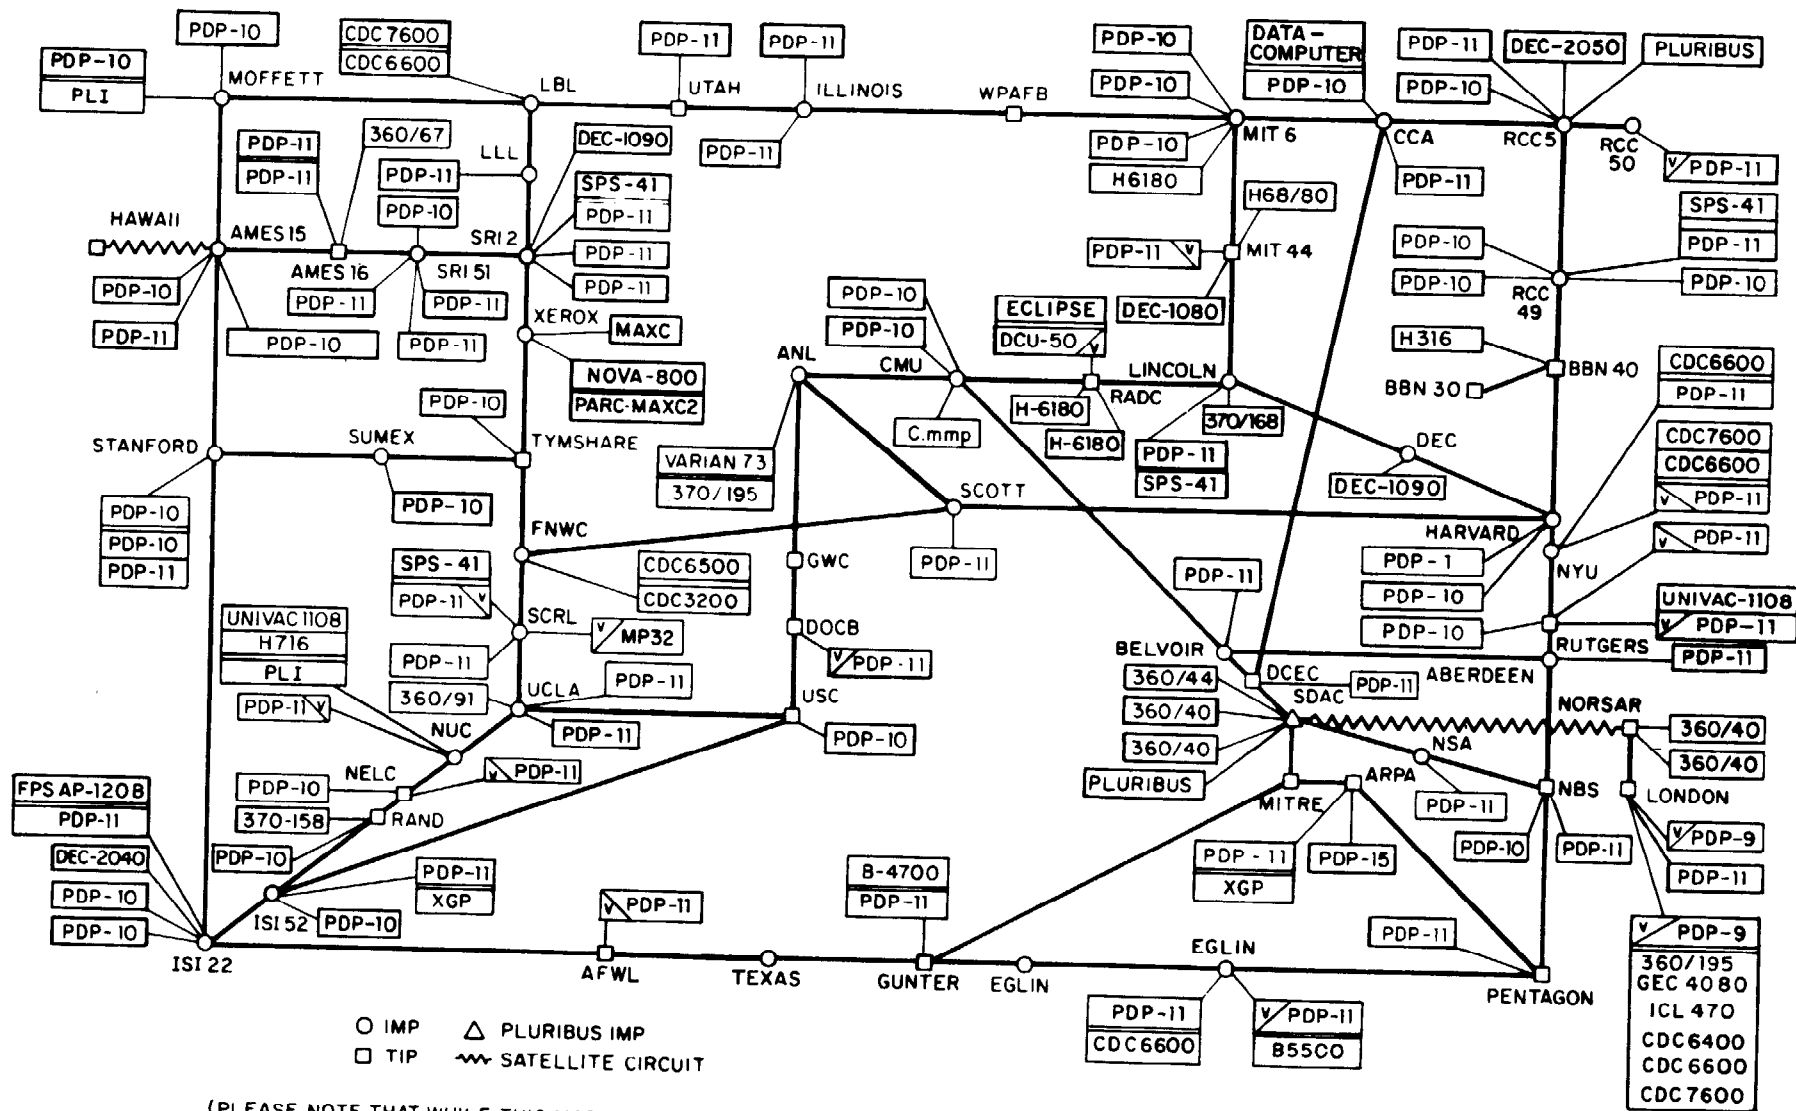


Figure 6.
ARPANET LOGICAL MAP, MARCH 1977



1.2.2.5 SYSTEM RELIABILITY AND BACKUP

System reliability has remained high over the past years; excellent under stable hardware and software conditions and degrading temporarily during debugging and development periods and during periods of difficult hardware problems. In general we take the system down for approximately 50 hours per month for scheduled hardware maintenance, file backup, and other maintenance. In addition we average from 10 to 15 hours per month in unscheduled downtime. During particularly difficult hardware or software difficulties we must absorb substantially more downtime.

1.2.2.6 PROGRAMMING LANGUAGES

Over the past years we or members of the SUMEX-AIM community have continued to maintain the major languages on the system at current release levels, have TENEXized several languages to improve efficiency, and have investigated a number of issues related to the efficiency of programs written in various LISP implementations and the exportability of programs. These issues are becoming increasingly critical in dealing with AI performance programs which have reached a level of maturity so that substantial, non-developmental user communities are growing. The following summarizes general accomplishments and the following section discusses in detail the work this past year in designing a machine-independent ALGOL-like system (MAINSAIL).

LISP Efficiency:

There has been an on-going debate among a number of projects over the best language to choose for developmental implementation of the various AI programs. The key issues include ease and flexibility of conceptual representation of program functions and objects, interactive debugging support, efficiency, and exportability. To date the predominant language choice for AIM research has been LISP and more particularly INTERLISP. These issues are important because they influence the time required to develop new AI programs and subsequently the incremental load placed on the SUMEX machine when in use. We recently attempted an evaluation of INTERLISP and ILISP including the relative efficiencies of the two languages and the level of assistance the language systems provide the user in developing programs. The tests were based on an implementation of a subset of REDUCE (a symbolic algebra manipulator). The results of several iterations in program refinement by experts in the respective languages were that the runtimes for the two versions were quite comparable (far less than the factor of 5-10 disparity predicted by ILISP enthusiasts). A more disquieting result was the substantial difference in runtimes depending on how particular functions were coded IN THE SAME LANGUAGE. It is apparent from the results that factors of 10 differences in time can result from a superficial implementation - expert programming insight is essential to efficient program performance. This is not a real surprise in that it is true of programming in any language - the problems may be increased by such a rich language as INTERLISP with such a wide array of

ways to do the same thing but with little guidance as to the relative costs. It has proven very difficult to quantify the "rules" for good programming. Mr. Masinter and Mr. Phil Jackson attempted to document good INTERLISP programming habits and issued a bulletin for SUMEX users.

A further impact of these data is that it is very difficult to simultaneously develop a new AI program and make the implementation highly efficient. With the iterations required to develop the conceptual design of the program, it is difficult to ensure its efficiency. This may lead to the need to reimplement the program after the basic development stabilizes to increase efficiency while still accommodating convenient and orderly further development. Such reimplementation may or may not be best done in LISP - this will depend on many factors including the nature of the program data structure requirements and anticipated further development efforts.

MAINSAIL Progress

SUMEX, in its role as a nationally shared computer resource, is an appropriate vehicle for the development of software unbound by the underlying machine environment. We have a built-in community of program developers acutely aware of the significance of providing their work to a broader base of users. This intersection of hardware capability, software expertise, and dedication to resource sharing presents a unique opportunity to promote a system designed for program sharing.

The MAINSAIL (3) project has three closely related goals:

- 1) Provide an integrated set of tools for the creation of efficient portable software on a variety of computer systems, and provide support and continued development of these tools in a form compatible across all implementations.
- 2) Study innovative approaches to portability, both hardware and software, and develop such approaches into effective tools.
- 3) Promote the development and distribution of portable software, advise and assist in its design, and evaluate its applicability.

By portable software we mean computer programs which may be executed on a variety of machines with few, if any, alterations. MAINSAIL itself will provide the initial example of portable software, since all of the system is written in the MAINSAIL language except for those parts which are determined by the host environment (hardware, instruction set, operating system, etc.). Even these parts are embedded within MAINSAIL.

(3) The MAINSAIL (MACHINE-INdependent SAIL) language is derived from SAIL, a programming language developed at Stanford University's Artificial Intelligence Laboratory. It is not compatible with SAIL, since SAIL was designed for a PDP-10 with TOPS-10, and hence contains machine-dependencies. However it has retained the basic attributes of SAIL as an extended ALGOL-like language. A summary of some of the features of the MAINSAIL language and their relationship to other languages is given in Appendix III on page 231 (see Book II).

There is a key distinction between MAINSAIL's approach to portability and the "classical" approach characterized by languages such as FORTRAN, ALGOL, LISP, COBOL and BASIC. These languages attempt to adhere to a single syntax standard which is separately implemented for each different computer system. Invariably these implementations have differences which preclude the creation of a program which is accepted by all. It is difficult, if not impossible, to define a language standard which is unambiguous and at the same time sufficiently comprehensible to provide the basis for compatible implementations. Furthermore, many implementors yield to the temptation to provide "enhancements" to the standard which immediately introduces machine and system dependencies.

MAINSAIL, on the other hand, provides a single system (written primarily in itself) which is employed at every site. This is made possible by its ability to compile itself into code for a variety of machines. Only the compiler's code generators and the runtime operating-system interfaces need be rewritten for each implementation. These parts of MAINSAIL are at a level which has already been defined by the machine-independent parts, and do not affect the language from the user's viewpoint. Thus the "language standard" has been reduced to a "semantic standard" which is surrounded by machine-independent software.

It remains to be seen whether the temptation to augment the language with machine-dependencies (for purposes of ultimate efficiency or to take advantage of particular local system features) can be overcome. Herein also lies the biggest "price" to be paid for exportability. The code emitted from the MAINSAIL compiler can be (and is, based on tests to date) at least as efficient as that from many machine-dependent compilers. On the other hand, special machine or operating system features that cannot be uniformly implemented may provide local optimizations at the cost of exportability or vice versa. We cannot effectively measure the extent of this cost at this stage.

DEVELOPMENT APPROACH

We do not underestimate the difficulty in obtaining the cooperation of a community which will span a wide variety of applications and hardware/software systems. If MAINSAIL is to obtain widespread use, it is crucial that it have an effective and credible base of support. The initial parts of MAINSAIL are just about ready for limited distribution. We want to maintain close supervision of this distribution, and insure that systems labelled as MAINSAIL are not altered without our approval. In this regard we are pursuing legal channels to safeguard the integrity of MAINSAIL software. We plan to take MAINSAIL through an orderly progression of development, and to avoid casual distribution with no provision for a solid base of maintenance and future growth.

REVIEW OF PROGRESS TO DATE

MAINSAIL has been under development for almost three years now. Beginning with an initial goal of converting the PDP-10 SAIL compiler to generate code for a PDP-11, several versions had been implemented on a PDP-10 and a PDP-11, and the groundwork had been laid for extending the system to a wider variety of machines. The current version was begun in August of 1976.

Early versions of MAINSAIL attempted to maintain close compatibility with the original SAIL, but in surveying a wider variety of machines (especially mini-computers), we concluded that this compatibility could be maintained only at the expense of portability. It was felt that MAINSAIL could contribute more by providing a truly portable system. Thus we began redesigning MAINSAIL, rebuilding from previous implementations. This effort has resulted in a new version which is still under development, and is now being tested on several systems.

Initial implementations of the current design are for DEC PDP-10's with the TENEX operating system and with the TOPS-10 operating system. The TENEX version is being tested at SUMEX and has been installed at one other TENEX site (Stanford - IMSSS). The TOPS-10 version was developed at SUMEX by using TENEX facilities which provide compatibility with TOPS-10. The Rutgers University PDP-10 facility was chosen for external testing since it is a standard TOPS-10 system, and can be accessed from SUMEX over a network. MAINSAIL is now undergoing preliminary testing there. A modified TOPS-10 version has been set up on the Stanford AI-lab's PDP-10, but also has not been open to general use.

Little additional work will be necessary to make the TENEX version execute on a DECSYSTEM-20 since TOPS-20 is derived from TENEX. However, some time will be needed to take full advantage of the extended instruction set of the KL-10. Two sites are available for TOPS-20 development: the LOTS facility at Stanford; and a machine at SRI, close to Stanford and accessible over a network. Both of these sites have expressed an interest in using MAINSAIL.

The PDP-11 has been chosen as the first mini-computer to be implemented. Code generators have been written for it but not debugged. Several variants of these code generators will be necessary to cover the full PDP-11 family.

MAINSAIL interfaces to three PDP-11 operating systems (RT-11, RSX-11 and UNIX) are now under development. All of these operating systems are available to the MAINSAIL project on PDP-11's at Stanford. RT-11 will be the first to be implemented. The mix of instruction sets, operating systems and configurations will be a good test of MAINSAIL's ability to provide a compatible implementation, even across this one family of computers. We expect the PDP-11 systems to be operational by this summer.

1.2.2.7 STANFORD AI HANDBOOK PROJECT

The AI Handbook is a compendium of short articles (3-5 pages each) about the projects, ideas, problems and techniques that make up the field of Artificial Intelligence. Over 150 articles have been drafted by researchers and students in the field, on topics ranging in depth from "Augmented Transaction Networks" (ATN's) to "An Overview of Natural Language Research", and covering the entire breadth of AI research: search, robotics, speech understanding, real-world applications, etc. An outline of the current contents of the handbook is given in Appendix II on page 225 (see Book II).

During the Spring of 1976 the final push for drafting new articles was completed, with some 60 articles produced by students during that quarter. Since then the process has begun of rewriting the various chapters of the Handbook to produce coherent manuscripts from the original work of five to ten authors. This effort involves rewriting articles for accuracy and completeness as well as integrating the 15 to 25 articles in a section into an editorially uniform and readable document. An editor has been added to the project team who will be responsible for maintaining a consistent format and style in the Handbook.

When completed, each chapter will be reviewed by experts in the appropriate research area before it is released to the public. At present, the chapter on Natural Language research is completed and being reviewed, and we expect that the sections on Search, Speech Understanding, Representation of Knowledge, and Automatic Programming will be completed during the next two months. During the Fall of 1977 the first seven chapters of the handbook will be published in preliminary form. Meanwhile, the handbook is already available to cooperative experts and critics on-line via the SUMEX-AIM network connections. We are considering maintaining the handbook on-line, with occasional hard-copy editions, and believe this method of "publication" may be a prototype for other encyclopedic monographs.

1.2.2.8 USER SOFTWARE AND INTRA-COMMUNITY COMMUNICATION

In addition to the system and language software development efforts of SUMEX, we have assembled or developed where necessary a broad range of utilities and user software. These include operational aids, statistics packages, DEC-supplied programs, improvements to the TOPS-10 emulator, text editors, text search programs, file space management programs, graphics support, a batch program execution monitor, text formatting and justification assistance, and magnetic tape conversion aids. We have also developed a number of user information assistance programs such as a "WHOIS" facility to recover names and affiliations of users and a "HELP" facility to locate on-line documentation of interest through key word searches.

Of major importance for our community effort is the set of tools for inter-user communications. We have enhanced the message sending and manipulation programs to better integrate text editing facilities for easier message preparation and reading. We have also developed a unique "bulletin board" system to deal with informal notes, thereby bridging a functional gap between formal system documents and private messages communications between individual users. The bulletin board system provides an informal and dynamic base for information about system facilities, lore, bugs, etc. or can provide a means for intra-project communication and coordination.

The system has been in operation for more than one year and has been exported to IMSSS (Stanford's other TENEX site) and USC-ECL. We have also proposed that the next generation of ARPANET information services provide for bulletin board-like facilities. At SUMEX-AIM there are 10 bulletin boards, 8 of which are project-specific. The main system bulletin board currently contains more than 140 bulletins under 85 topics covering system status announcements,

explanations of recent crashes, hardware troubles and monitor upgrades, new developments, bugs, and little-documented features of our programming languages and utilities. Project bulletin boards have been used for notices and minutes of meetings, references to and abstracts of papers, coordination of on-going developments, vacation schedules, documentation and announcements of various kinds.

Current Bulletin Board features include:

Multiple bulletin boards (public, private, general, specific, etc.).

Topics and subtopics (separated by periods) may be nested to any depth.

Expire dates for each bulletin, after which they are removed automatically.

Interest-list-of-topics for each user allows him to be notified about new bulletins he is interested in and to ignore others.

Users notified when new bulletins arrive, by running BBCHECK (the bulletin-board MAIL CHECK) or by mail.

Help and browsing facilitated in a variety of ways (? can be typed anywhere, general and command-specific help provided).

Command structure modelled after the TENEX EXEC, with conscious attention to human-engineering.

Companion program BBREAD is a bulletin-board READMAIL.

Companion program BBNEWS types out a directory listing of any new bulletins.

1.2.2.9 DOCUMENTATION AND EDUCATION

We have spent considerable effort to develop, maintain, and facilitate access to our documentation so as to accurately reflect available software. The HELP and Bulletin Board systems have been important in this effort. We have limited manpower for user assistance. In general, users are responsible for their own software development and maintenance. The SUMEX staff, however, (including Lederberg and Rindfleisch) share the responsibilities for system level assistance to users, tracking down bugs, reviewing user suggestions, etc. The terminal linking facilities of TENEX have been valuable tools to assist remote user groups and also for system users to communicate with each other. With the recent initial release of the MAINSAIL system on selected machines, we are becoming increasingly involved in describing MAINSAIL and advising user projects in its possible applications.

1.2.2.10 SOFTWARE COMPATIBILITY AND SHARING

At SUMEX-AIM we firmly believe in importing rather than reinventing software where possible. At SUMEX many avenues exist for sharing between the system staff, various user projects, other facilities, and vendors. In the past

without communication networks, the system vendor served as the focal point for distribution of most software to user sites. Since the process of distributing tapes (and particularly of handling bug reports and user suggestions) was very slow, it was common for sites to take a version of a program and then modify and maintain it locally. This caused a proliferation of home-grown versions of software. Similar impediments have existed to the dissemination of user software. User organizations like SHARE and DECUS have helped to overcome these problems but communication is still cumbersome. The advent of fast and convenient communication facilities coupling communities of computer facilities has the potential of making a major difference in facilitating inter-group cooperation and to lower these barriers.

The TENEX sites on the ARPANET have been interacting increasingly with each other to develop new software systems. This functions effectively to build communication around the network and promote a functional division of labor and expertise. The other major advantage is that as a by-product of the constant communication about particular software, personal connections between staff members of the various sites develop. These connections serve to pass general information about software tools and to encourage the exchange of ideas among the sites. Certain common problems are now regularly discussed on a multi-site level. We continue to draw significant amounts of system software from other ARPANET sites, reciprocating with our own local developments. Interactions have included mutual backup support, hardware configuration experiments, operating system enhancements, utility or language software, and user project collaborations. We have been able to import many new pieces of software and improvements to existing ones in this way. Examples of imported software include the message manipulation program MSG, TENEX SAIL, TENEX SOS, INTERLISP, the RECORD program, ARPANET host tables, and many others. Reciprocally, we have exported our contributions such as the drum page migration system, KI-10 page table efficiency improvements, GTJFN enhancements, PUB macro files, the bulletin board system, SNDMSG enhancements, our BATCH monitor, etc. The most recent example of this cooperative use of networks is in the preliminary export of MAINSAIL.

1.2.2.11 RESOURCE MANAGEMENT

ORGANIZATION AND PROCEDURES

The SUMEX-AIM resource is administered within the Genetics Department of the Stanford University Medical School, Professor Lederberg's "main office", though he also holds appointments in the Computer Science Dept. and the Human Biology program. Its mission, locally and nationally, entails both the recruitment of appropriate research projects interested in medical AI applications and the catalysis of interactions among these groups and the broader medical community. User projects are separately funded and autonomous in their management. They are selected for access to SUMEX on the basis of their scientific and medical merits as well as their commitment to the community goals of SUMEX. Currently active projects span a broad range of application areas such as clinical diagnostic consultation, molecular biochemistry, belief systems modeling, mental function modeling, and instrument data interpretation (see Section 6 on page 41 in Book II).

EXECUTIVE AND ADVISORY COMMITTEE ORGANIZATION

As the SUMEX-AIM project is a multilateral undertaking by its very nature, we have created several management committees to assist in administering the various portions of the SUMEX resource. As defined in the SUMEX-AIM management plan adopted at the time the initial resource grant was awarded, the available facility capacity is allocated 40% to Stanford Medical School projects, 40% to national projects, and 20% to common system development and related functions. Within the Stanford aliquot, Dr. Lederberg has established an advisory committee to assist him in selecting and allocating resources among projects appropriate to the SUMEX mission. The current membership of this committee is listed in Appendix V (see Book II).

For the national community, two committees serve complementary functions. An Executive Committee oversees the operations of the resource as related to national users and makes the final decisions on authorizing admission for projects. It also establishes policies for resource allocation and approves plans for resource development and augmentation within the national portion of SUMEX (e.g., hardware upgrades, MAINSAIL development priorities, etc.). The Executive Committee oversees the planning and implementation of the AIM Workshop series currently implemented under Prof. S. Amarel of Rutgers University and assures coordination with other AIM activities as well. The committee will play a key role in assessing the possible need for additional future AIM community computing resources and in deciding the optimal placement and management of such facilities. The current membership of the Executive committee is listed in Appendix V (see Book II).

Reporting to the Executive Committee, an Advisory Group represents the interests of medical and computer science research relevant to AIM goals. The Advisory Group serves several functions in advising the Executive Committee; 1) recruiting appropriate medical/computer science projects, 2) reviewing and recommending priorities for allocation of resource capacity to specific projects based on scientific quality and medical relevance, and 3) recommending policies and development goals for the resource. The current Advisory Group membership is given in Appendix V (see Book II).

These committees have actively functioned in support of the resource. Except for the meetings held during the AIM workshops, the committees have met by telephone conference owing to the size of the groups and to save the time and expense of personal travel to meet face to face. These telephone meetings, in conjunction with terminal access to related text materials, have served quite well in accomplishing the agenda business and facilitate greatly the arrangement of meetings. Other solicitations of advice requiring review of sizable written proposals are done by mail.

We will continue to work with the management committees to recruit the additional high quality projects which can be accommodated and to evolve resource allocation policies which appropriately reflect assigned priorities and project needs. We hope to make more generally available information about the various projects both inside and outside of the community and thereby to promote the kinds of exchanges exemplified earlier and made possible by network facilities.

NEW PROJECT RECRUITING

The SUMEX-AIM resource has been announced through a variety of media as well as by correspondence, contacts of NIH-BRP with a variety of prospective grantees who use computers, and contacts by our own staff and committee members. The number of formal projects that have been admitted to SUMEX has more than doubled since the start of the project; others are working tentatively as pilot projects or are under review.

We have prepared a variety of materials for the new user ranging from general information such as is contained in a brochure (see Appendix VI in Book II) to more detailed information and guidelines for determining whether a user project is appropriate for the SUMEX-AIM resource. Dr. E. Levinthal has prepared a questionnaire to assist users seriously considering applying for access to SUMEX-AIM (see Appendix VII in Book II). Pilot project categories have been established both within the Stanford and national aliquots of the facility capacity to assist and encourage projects just formulating possible AIM proposals pending their application for funding support and in parallel formal application for access to SUMEX. Pilot projects are approved for access for limited periods of time after preliminary review by the Stanford or AIM Advisory Group as appropriate to the origin of the project.

These contacts have sometimes done much more than provide support for already-formulated programs. For example, Prof. Feigenbaum's group at Stanford has initiated a major collaborative effort with Dr. Osborn's group at the Institutes of Medical Sciences in San Francisco. This project in "Pulmonary Function Monitoring and Ventilator Management - PUFF/VM" (see Section 6.4.6 on page 197 in Book II) originated as a pilot request to use MLAB in a small way for modeling. Subsequently the AI potentialities of this domain were recognized by Feigenbaum, Nii, and Osborn who have submitted a joint proposal to NIH and have a pilot status at present.

The following lists the fully authorized projects currently comprising the SUMEX-AIM community (see Section 6 in Book II for more detailed descriptions). The nucleus of five projects that were authorized at the initial funding of the resource in December 1973 are marked by "<*>".

National Community -

- 1) Acquisition of Cognitive Procedures (ACT); Dr. J. Anderson (Yale University)
- <*> 2) Higher Mental Functions Project; K. Colby, M.D. (University of California at Los Angeles)
- 3) INTERNIST Project; J. Myers, M.D. and Dr. H. Pople (University of Pittsburgh)
- 4) Medical Information Systems Laboratory (MISL); J. Wilensky, M.D. and Dr. B. McCormick (University of Illinois at Chicago Circle)
- <*> 5) Rutgers Computers in Biomedicine; Dr. S. Amarel (Rutgers University)
- 6) Chemical Synthesis Project (SECS); Dr. T. Wipke (University of California at Santa Cruz)

Stanford Community -

- <*> 1) DENDRAL Project; Drs. C. Djerassi, J. Lederberg, and E. Feigenbaum
- 2) Large Multi-processor Arrays (HYDROID); Dr. G. Wiederhold
- 3) Molecular Genetics Project (MOLGEN); Drs. J. Lederberg, E. Feigenbaum, and N. Martin
- <*> 4) MYCIN Project; S. Cohen, M.D. and Dr. B. Buchanan
- <*> 5) Protein Structure Modelling; Drs. J. Kraut and S. Freer (University of California at San Diego) and E. Feigenbaum (Stanford)

As an additional aid to new projects or collaborators with existing projects, we provide a limited amount of funds for use to support terminals and communications needs of users without access to such equipment. We are currently leasing 6 terminals and 4 modems for users as well as 4 foreign exchange lines to better couple the Rutgers project into the TYMNET and a leased line between Stanford and U. C. Santa Cruz for the Chemical Synthesis project.

STANFORD COMMUNITY BUILDING

The Stanford community has undertaken several internal efforts to encourage interactions and sharing between the projects centered here. Professor Feigenbaum organized a seminar class with the goal of assembling a handbook of AI concepts, techniques, and current state-of-the-art. This project has had enthusiastic support from the students and substantial progress made in preparing many sections of the handbook as reported earlier. An outline of the material being prepared can be found in Appendix II on page 225 (see Book II). Several examples of completed articles are given in Appendix I on page 202 (see Book II).

A second community-building effort was a mini-conference on AI held at Stanford in January 1976. This 3 day series of meetings featured presentations by each of the local projects and comparative discussions of approaches to current problems in AI research such as knowledge representations, production system strategies and rule formation, etc. Weekly informal lunch meetings (SIGLUNCH) are also held between community members to discuss general AI topics, concerns and progress of individual projects, or system problems as appropriate as well as having a number of outside invited speakers.

AIM WORKSHOP SUPPORT

The Rutgers Computers in Biomedicine resource (under Dr. Saul Amarel) has organized a series of workshops devoted to a range of topics related to artificial intelligence research, medical needs, and resource sharing policies within NIH. Meetings have been held for the past two years at Rutgers and another is planned for this summer. The SUMEX facility has acted as a prime computing base for the workshop demonstrations. We expect to continue this support for future workshops. The AIM workshops provide much useful information about the strengths and weaknesses of the performance programs both in terms of criticisms from other AI projects and in terms of the needs of practicing medical people. We plan to continue to use this experience to guide the community building aspects of SUMEX-AIM.

RESOURCE ALLOCATION POLICIES

As the SUMEX facility has become increasingly loaded, a number of diverse and conflicting demands have arisen which require controlled allocation of critical facility resources (file space and central processor time). We have already spelled out a policy for file space management; an allocation of file storage is defined for each authorized project in conjunction with the management committees. This allocation is divided among project members in any way desired by the individual principal investigators. System allocation enforcement is implemented by project each week. As the weekly file dump is done, if the aggregate space in use by a project is over its allocation, files are archived from user directories over allocation until the project is within its allocation.

We have recently implemented system scheduling controls to attempt to maintain the 40:40:20 balance in terms of CPU utilization (see page 14). The initial complement of user projects justifying the SUMEX resource was centered to a large extent at Stanford. Over the first term of the SUMEX grant, a substantial growth in the number of national projects was realized. During the same time the Stanford group of projects has matured as well and in practice the 40:40 split between Stanford and non-Stanford projects is not ideally realized (see Figure 8 on page 38 and the tables of recent project usage on page 40). Our job scheduling controls bias the allocation of CPU time based on percent time consumed relative to the time allocated over the 40:40:20 community split. The controls are "soft" however in that they do not waste computer cycles if users below their allocated percentages are not on the system to consume the cycles. The operating disparity in CPU use to date reflects a substantial difference in demand between the Stanford community and the developing national projects, rather than inequity of access. For example, the Stanford utilization is spread

over a large part of the 24-hour cycle, while national-AIM users tend to be more sensitive to local prime-time constraints. (The 3-hour time-zone phase shift across the continent is of substantial help in load-balancing.) For the present, we propose to continue our policy of "soft" allocation enforcement for the fair split of resource capacity. If necessary to assure proper apportionment, we can implement a pie-slice reservation system to more rigidly control the allocations.

Our system also categorizes users in terms of access privileges. These comprise fully authorized users, pilot projects, guests, and network visitors in descending order of system capabilities. We want to encourage bona fide medical and health research people to experiment with the various programs available with a minimum of red tape while not allowing unauthenticated users to bypass the advisory group screening procedures by coming on as guests. So far we have had relatively little abuse compared to what other network sites have experienced, perhaps on account of the personal attention that senior staff gives to the logon records, and to other security measures. However, the experience of most other computer managers behooves us to be cautious about being as wide-open as might be preferred for informal service to pilot efforts and demonstrations. We will continue developing this mechanism in conjunction with management committee policy decisions.

1.2.2.12 SUMMARY OF RESOURCE USAGE

The following data give an overview of SUMEX-AIM resource usage. There are five sub-sections containing data respectively for 1) monthly CPU time consumed, 2) resource usage by community (AIM and Stanford), 3) resource usage by project, 4) recent diurnal loading data, and 5) Network usage data.


MONTHLY CPU TIME CONSUMED

Figure 7. Monthly CPU Time Consumed

RELATIVE SYSTEM LOADING BY COMMUNITY

The SUMEX resource is divided, for administrative purposes, into 3 major communities: user projects based at the Stanford Medical School, user projects based outside of Stanford (national AIM projects), and common systems development efforts. As defined in the resource management plan approved by BRP at the start of the project, the available resource in terms of CPU capacity and file space will be divided between these communities as follows:

Stanford	40%
AIM	40%
Staff	20%

The "available" resources to be divided up in this way are those remaining after various monitor and community-wide functions are accounted for. These include such things as job scheduling, overhead, network service, file space for subsystems and documentation, etc.

The monthly usage of CPU and file space resources for each of these three communities relative to their respective aliquots is shown in the plots in Figure 8 and Figure 9. It is clear that the Stanford projects have held an edge in system usage despite our efforts at resource allocation and the substantial voluntary efforts by the Stanford community to utilize non-prime hours. This reflects the development of the Stanford group of projects relative to those getting started on the national side and has correspondingly accounted for much of the progress in AI program development to date.




Figure 8. CPU Usage by Community




Figure 9. File Space Usage by Community

INDIVIDUAL PROJECT AND COMMUNITY USAGE

The table following shows cumulative resource usage by project in the past grant year. The data displayed include a description of the operational funding sources (outside of SUMEX-supplied computing resources) for currently active projects, total CPU consumption by project (Hours), total terminal connect time by project (Hours), and average file space in use by project (Pages, 1 page = 512 computer words). These data were accumulated for each project for the months between May 1976 and April 1977. Again the well developed use of the resource by the Stanford community can be seen. It should be noted that the Stanford projects have voluntarily shifted a substantial part of their development work to non-prime time hours which is not shown in these cumulative data. It should also be noted that a significant part of the DENDRAL and MYCIN efforts, here charged to the Stanford aliquot, support development efforts dedicated to national community access to these systems. The actual demonstration and use of these programs by extramural users is charged to the national community in the "AIM USERS" category, however.

RESOURCE USE BY INDIVIDUAL PROJECT

STANFORD COMMUNITY	CPU (Hours)	CONNECT (Hours)	FILE SPACE (Pages)
1) DENDRAL PROJECT "Resource Related Research Computers and Chemistry" NIH RR-00612-08 (3 yrs. 1977-80) ARPA DAHC-15-73-C-0435 (2 yrs. 1977-79)	1181.64	19657.56	13058
2) HYDROID PROJECT "Distributed Processing and Problem Solving" ARPA DAHC-15-73-C-0435	40.92	924.49	239
3) MOLGEN PROJECT NSF MCS76-11649 NSF MCS76-11935 (2 yrs. 1976-78)	85.61	2487.73	1853
4) MYCIN PROJECT "Computer-based Consult. in Clin. Therapeutics" HEW HS-01544 (2 yrs. 1977-79) NSF (2 yrs. 1977-79)	410.87	6640.75	6688
5) PROTEIN STRUCT MODELING "Heuristic Comp. Applied to Prot. Crystallog." NSF DCR 74-23461 (2 yrs. 1977-79) ARPA DAHC 15-73-C-0435	159.80	2894.19	2477
6) AIHANDBOOK PROJECT	26.46	464.42	639
7) PILOT PROJECTS (see reports in Section 6.3 in Book II)	327.67	5919.33	3506
	-----	-----	-----
COMMUNITY TOTALS	2232.97	38988.47	28460

NATIONAL AIM COMMUNITY

1) ACT PROJECT "Acquisition of Cognitive Procedures" NIMH MH29353 ONR N0014-77-6-0242	57.02	1195.84	986
2) HIGHER MENTAL FUNCTIONS "Computer Models in Psychiatry and Psychother." NIH MH-27132-02 (2 yrs.) UCLA NPI Gen. Res.	206.03	2680.16	2198
3) INTERNIST PROJECT (DIALOG) "Computer Model of Diagnostic Logic" BHRD MB-00144-03 (3 yrs.)	205.20	2721.26	3535
4) MISL PROJECT "Medical Information Systems Laboratory" US-PHS-MB00114-03 (3 yrs.)	9.27	380.05	876
5) RUTGERS PROJECT "Computers in Biomedicine" NIH RR-00643-05 (3 yrs.)	139.63	2433.43	10862
6) SECS PROJECT "Chemical Synthesis"	308.96	4374.03	4515
7) AIM PILOT PROJECTS (see reports in Section 6.4 in Book II)	40.91	1326.56	1558
8) AIM Administration	11.13	383.22	1762
9) AIM Users	56.89	672.35	362
	-----	-----	-----
COMMUNITY TOTALS	1035.04	16166.90	26654

SUMEX STAFF AND SYSTEM

1) Staff	903.07	23198.86	11919
2) Miscellaneous	80.87	2508.98	1721
3) Operations	1505.50	63113.94	32382
	-----	-----	-----
COMMUNITY TOTALS	2489.44	88821.78	46022
	=====	=====	=====
RESOURCE TOTALS	5757.45	143977.15	101136

SYSTEM DIURNAL LOADING VARIATIONS

The following figures give a picture of the recent variations in diurnal SUMEX system load, taken during March 1977. The plots include:

- Figure 10 - Total number of jobs logged in to the system
- Figure 11 - Percent of total CPU time used by logged in jobs (maximum is 200% for dual processor capacity)
- Figure 12 - Percent of total CPU time consumed as overhead; I/O wait, core management, scheduling, etc. (maximum = 200%)
- Figure 13 - Balance set size (number of jobs in core)
- Figure 14 - Number of runnable jobs (whether or not in core)

The abscissa for these plots is broken into 20 minute intervals throughout the day. The ordinate for each interval is the average of all the daily measurements for that interval over the weekdays during March 1977. A daily measurement for a given 20 minute interval is in turn an average of the appropriate statistic sampled every 10 seconds. Since these plots display overall average data, they give representative illustration of the general characteristics of diurnal loading. There are, of course, substantial fluctuations in the quantities measured from day to day as well and for some, also on time scales shorter than the intervals displayed in the figures. For example in Figure 14, the number of runnable jobs (equivalent to the system "load average") shows a fairly smooth curve peaking at 6.7 jobs. On both a scale of minutes and from day to day, however, the number of runnable jobs will vary from only a few to 12 or more. This fluctuation is not shown in these average plots but also plays a role in the responsiveness of the system.

In the heading of each plot are shown range statistics for the measurement over various parts of the day. Range data include the minimum value "Low", average value "Ave", and maximum value "High". The first line of the heading gives the range over the whole day and on succeeding lines, "Prime Time" covers 6:00-18:00 Pacific time and "Non Prime Time" covers the remaining night time hours.

It can be noted in Figure 12 that the current overhead level for the dual processor system is quite high (about 33% per processor). This is because of the limited memory size (256K words) we currently have and the resulting increase in swapping interrupt rate and I/O wait time. We have a proposal pending with the AIM Executive committee to augment our memory which should reduce this overhead down to our earlier single processor levels (about 15-20% per processor).

Figure 10. Average Diurnal Loading (3/77): Total Number of Jobs




Figure 11. Average Diurnal Loading (3/77): Percent Time Used




Figure 12. Average Diurnal Loading (3/77): Percent Overhead




Figure 13. Average Diurnal Loading (3/77): Balance Set - Jobs in Core





Figure 14. Average Diurnal Loading (3/77): Runnable Jobs



1.2.2.13 NETWORK USAGE STATISTICS

NETWORK USAGE PLOTS

The plots in Figure 15 show the major billing components for SUMEX-AIM TYMNET usage. These include the total connect time for terminals coming into SUMEX and the total number of characters transmitted over the net. The ratio of characters received at SUMEX to characters sent to the terminal is about 1:12 over our period of usage. Also shown for recent months is a plot of ARPANET connect time which tracks the corresponding data for TYMNET usage fairly closely. No data for "character" transmission is available for ARPANET since file transfers and terminal traffic use different byte sizes and these data are not resolved and maintained for the ARPANET.





Figure 15. TYMNET and ARPANET Usage Data

1.2.2.14 PUBLICATIONS

The following are publications for the SUMEX staff and have included papers describing the SUMEX-AIM resource and on-going research as well as documentation of system and program developments. Publications for individual collaborating projects are detailed in their respective reports (see Section 6 on page 41 in Book II).

- [1] Carhart, R.E., Johnson, S.M., Smith, D.H., Buchanan, B.G., Dromey, R.G., and Lederberg, J, "Networking and a Collaborative Research Community: a Case Study Using the DENDRAL Programs", ACS Symposium Series, Number 19, COMPUTER NETWORKING AND CHEMISTRY, Peter Lykos (Editor), 1975.
- [2] Levinthal, E.C., Carhart, R.E., Johnson, S.M., and Lederberg, J., "When Computers Talk to Computers", Industrial Research, November 1975
- [3] Wilcox, C. R., "MAINSAIL - A Machine-Independent Programming System," Proceedings of the DEC Users Society, Vol 2, No 4, Spring 1976.

Mr. Clark Wilcox also chaired the session on "Languages for Portability" at the DECUS DECsystem10 Spring '76 Symposium.

In addition as reported earlier, a substantial effort has gone into developing, upgrading, and extending documentation about the SUMEX-AIM resource, the SUMEX-TENEX system, the many subsystems available to users, and MAINSAIL. These efforts include a number of major documents (such as SOS, PUB, and TENEX-SAIL manuals) as well as a much larger number of document upgrades, user information and introductory notes, an ARPANET Resource Handbook entry, and policy guidelines (see Appendix VI, and Appendix VII in Book II). Publications for individual user projects are summarized in the respective reports (see Section 6 in Book II).

BOOK II

Collaborating Project Reports and Supporting Appendixes

The following sections detail the reports and plans for SUMEX-AIM collaborating projects and also contain additional information in the form of appendixes relating to the core resource progress and operation. The heading and page numbering of these sections does not continue sequentially from that of the Book I progress report. The discontinuity reflects the initial organization of this material as part of our renewal grant application.

6 COLLABORATIVE PROJECT PROGRESS AND OBJECTIVES

The following subsections report on the collaborative use of the SUMEX facility including the formally authorized projects within the Stanford and AIM aliquots and the various "pilot" efforts currently under way. These project descriptions and comments are the result of a solicitation for contributions sent to each of the project Principal Investigators requesting the following information:

- I) Summary of research program
 - A) Technical goals
 - B) Medical relevance and collaboration
 - C) Progress summary
 - D) Up-to-date list of publications
- II) Interactions with the SUMEX-AIM resource
 - A) Examples of collaborations and medical use of programs via SUMEX
 - B) Examples of sharing, contacts and cross-fertilization with other SUMEX-AIM projects (via workshops, system facilities, personal contact, etc.)

We believe that the reports of the individual projects speak for themselves as rationales for participation; in any case the reports are recorded as submitted and are the responsibility of the indicated project leaders.

6.1 STANFORD PROJECTS

The following group of projects is formally approved for access to the Stanford aliquot of the SUMEX-AIM resource. Their access is based on review by the Stanford Advisory Group and approval by Professor Lederberg as Principal Investigator. As noted previously, the DENDRAL project was the historical core application of SUMEX. Although this is described as a "Stanford project," a significant part of the development effort and of the computer usage is dedicated to national collaborator-users of the DENDRAL programs.

6.1.1 DENDRAL PROJECT

DENDRAL - Resource Related Research - Computers & Chemistry

Carl Djerassi, Principal Investigator
Professor of Chemistry
Stanford University

I. OVERVIEW OF RESEARCH ACTIVITIES

Technical Goals

Our research, development and future plans focus on both the question of structure elucidation in general and the problem of providing computer assistance to scientists engaged in specific aspects of this important activity.

A simplified representation of major milestones in solving unknown biomolecular structures by manual methods is presented in Figure 1.

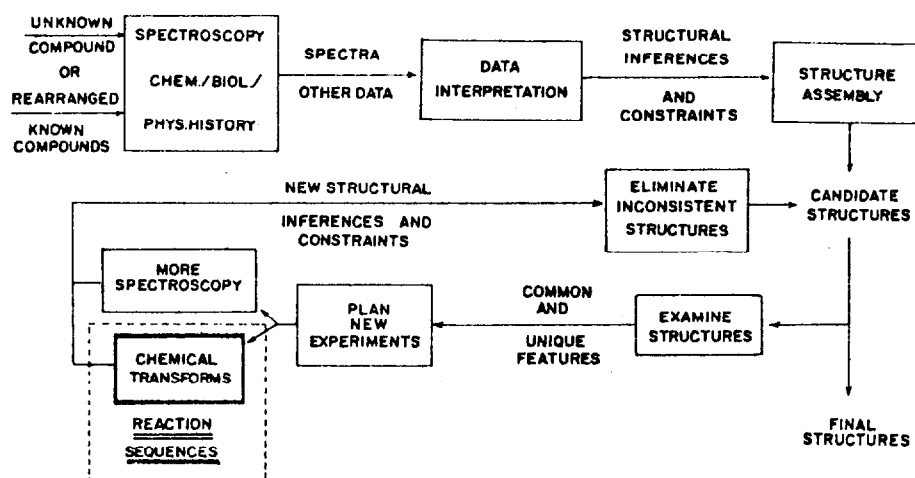


Figure 1. Important steps in manual solution of structures of unknown chemical compounds.

These steps, indicated as separate boxes, may be performed explicitly or implicitly. There are considerably more complex relationships among the boxes of Fig. 1 than are indicated when structures are actually solved. Nevertheless, the Figure provides a good introduction to both our recent work and our future directions. We describe briefly each of the milestones in the following paragraphs. More detailed discussions of each topic follow in subsequent sections.

The first step in identification of an unknown structure is to separate it from other components in a potentially complex mixture and to isolate it in reasonably pure form. These steps are performed by scientists, frequently with the assistance of various instruments. Although our research is not directed toward any part of this separation and isolation procedure (except insofar as these procedures also yield data which are subject to computer-assisted interpretation), information about the chemical and physical characteristics of the compound may be crucial to further efforts to determine its structure.

Depending on the quantity of sample available and its characteristics, various spectroscopic and additional chemical data are then collected on the unknown. A mass spectrum is frequently obtained, e.g., from a combined gas chromatograph/mass spectrometer (GC/MS) system. An important part of our recent proposal to the NIH is directed toward automation of combined GC/MS systems operated at high mass spectrometer resolving powers. Data on elemental compositions and relative ion abundances are then available in computer-readable form for further analysis (see MSRANK). The chemist possess an armamentarium of spectroscopic techniques which can be brought to bear on a structure. One advantage of our work is that any data so obtained can be used to help solve the structure as long as it can be expressed, manually or by computer, in substructural statements about the unknown.

The next important phase in structure elucidation is interpretation of the available data (Fig. 1) in terms of structural features of the molecule. These interpretations may be in terms of known structural units ("superatoms", polyatomic aggregates of atoms in known configurations), or in terms of structural units, ring sizes, proton or carbon distributions. The latter set of features represents constraints on the kinds of structures which are possible. Our efforts in the area of computer-assisted data interpretation are focussed on mass spectral and carbon-13 nuclear magnetic resonance (¹³CMR) data. We are developing general approaches to automated analysis of these data in terms of structural features of unknowns.

Our recent efforts are summarized in Figure 2, and discussed in detail subsequently. We have been concerned with use of these data from two points of view, planning and prediction (Fig. 2). During planning, experimental data are examined in order to extract specific structural information to be used in assembling candidate structures. In prediction each candidate structure is tested to determine how closely its predicted spectrum agrees with the observed spectrum. The candidates can be ranked accordingly. The Meta-DENDRAL research is directed toward determination of rules of spectroscopic data which can be used either for planning or prediction (see below).

Given possible structural fragments of the complete molecule and constraints on how these fragments may be assembled into complete molecules, a process of structural assembly follows (Fig. 1). There has been no proven algorithm for solving this problem prior to earlier work supported by the current grant. Traditionally, this process has been left to manual, pencil and paper work. Our CONGEN program, which was designed to solve this problem, is the farthest advanced of programs designed to assist in various aspects of structure elucidation. It performs the structural assembly process, under constraints, and

DATA INTERPRETATION"PLANNING"

EXTRACTION OF STRUCTURAL
INFORMATION DIRECTLY FROM
SPECTROSCOPIC DATA.

1. MASS SPECTRA - MDGGEN
2. ¹³CNMR

PREDICTION

USE OF SPECTROSCOPIC
DATA TO RANK
CANDIDATE STRUCTURES.

1. MSPRUNE, MSPRED
2. ¹³CNMR

← ↗
META - DENDRAL

FORMATION OF RULES TO BE
USED FOR BOTH PLANNING
AND PREDICTION.

Figure 2. Relationship between use of rules in either planning or prediction.
Both approaches are used in utilizing data for structure elucidation.

allows the scientist using the program to examine structural candidates and remove those deemed implausible (Fig. 1). A large portion of our recent and future work is directed toward improving the CONGEN program and building other facilities around it (see later sections). We have demonstrated the utility of CONGEN in structural studies, and subsequent sections discuss our recent developments and applications of CONGEN as well as our interactions with other scientists desiring access to our programs.

Given a set of structural candidates, the experimenter examines them to determine what experiments might be performed to focus on the correct structure by stepwise rejection of alternative hypotheses. When there are only a small number of possibilities under consideration, manual methods suffice. But CONGEN provides the capability for exhaustive enumeration of structural possibilities at a point in a structural problem when there may be many hundreds of possibilities. It is very difficult to examine these structures and plan experiments by hand. We have begun exploring ways to provide computer assistance to this important aspect of structure elucidation. We refer to this research area as the Experiment Planner, discussed in more detail below.

When new experiments have been planned the researcher carries them out and uses the results as additional constraints on the structural candidates (Fig. 1). New experiments may include collecting of additional spectroscopic data or performing a sequence of chemical reactions on the unknown. The latter experiments may be chosen to convert the unknown into a related compound which possesses physical or chemical properties more amenable to analysis. During the past year we have developed a program to assist scientists in carrying out representations of chemical reactions in the computer and eliminating undesired structural candidates based on constraints exercised on the products of the reaction. This work is described in two subsequent sections. One section describes use of the program, which we call REACT, to explore structural possibilities exactly as outlined above. A later section describes recent progress in increasing the power of REACT.

Medical Relevance

Structure elucidation is a fundamental problem for medical practice and biomedical research. For example, we are collaborating with physicians in the Department of Pediatrics who monitor the body fluids of newborn infants in order to detect abnormal compounds. Much of the research leading to new drugs and new methods for synthesizing drugs also depends on careful analysis and identification of molecular structures of compounds. The computer tools that we are developing will aid in the determination of molecular structures by giving working scientists help with data collection, data interpretation, hypothesis testing and, most important, systematic consideration of all molecular structures that are consistent with the interpretations of the available data.

PROGRESS SUMMARY

Experiment Planner

We have begun preliminary considerations of design and implementation of an experiment planner. This program will assist chemists in designing the most effective set of experiments to perform to solve the structure. Although the experiment planner will be a future activity of our group, we are developing and using other structure manipulation functions which will provide groundwork for future developments.

One important aspect of experiment planning is the ability to examine in some way the set of candidate structures. Although many can be drawn for visual review, drawing is impractical when dozens or hundreds of structures are involved. To assist persons using CONGEN in reviewing their structures we have developed a function auxiliary to CONGEN which we call SURVEY.

SURVEY

FUNCTION: AIDS IN PERCEPTION OF ANY OF A
PRE-SPECIFIED SET OF STRUCTURAL
FEATURES IN A GROUP OF
STRUCTURAL CANDIDATES.

E.G. A) FUNCTIONAL GROUPS
B) TERPENOID SKELETONS
C) AMINO ACID SKELETONS

Figure 3. Function of the SURVEY program and examples of recent application areas.

The function of SURVEY is summarized in Figure 3. SURVEY simply acts as a reminder to the scientist of the presence or absence of certain structures or structural features. During the past year we have used SURVEY extensively. For example, we have used it to detect implausible functional groups in a set of candidate structures, using a file of substructures representing a wide variety of functionalities. In many problems, implausible functional groups are forgotten and CONGEN is never constrained to remove them. Another example of use of SURVEY is in conjunction with collaborative work with persons in the

Department of Genetics. In analysis of serum or urinary metabolites in patients of high risk of metabolic disorder, we have had occasion to use CONGEN in exploration of unknown structures [Report HPP-77-11]. Some of these structures could formally be conjugates of amino acids with organic acids. If so, such structures will possess backbones of naturally-occurring amino acids. SURVEY was used to provide a summary of which structural candidates possessed such amino acid skeletons.

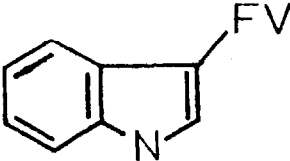
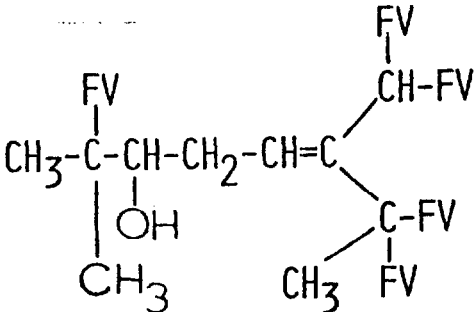
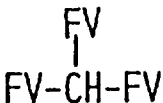
We have recently used SURVEY in a related application involving the structure of "polyathenol", discussed by LeBoeuf, et al. (Figure 4). Superatoms and constraints supplied to CONGEN to derive structural candidates are summarized in Fig. 4.

We summarize in Figure 5 the structural possibilities which resulted. There are five structures possessing a bicyclo[2.1.1] system, and six which possess a bicyclo[4.3.1] system (Fig. 5, top). These structures are energetically less favorable. For example, several possess a double bond at a bridgehead atom, which violates Bredt's Rule. There remain, however, 11 structures which are not formally excluded by data presented by LeBoeuf, et al. Because these workers based their structural assignment on biogenetic grounds, we used SURVEY and REACT to test their hypothesis. We have, in computer-accessible libraries, known terpenoid ring systems which can be used within SURVEY to test sets of structures for known skeletons. None of the 22 structural candidates possesses a previously known skeleton. Because the authors postulated a relationship to a known skeleton via a single methyl shift, we used REACT to exercise a single methyl shift in all possible ways on each of the 22 candidates. SURVEY was then used to test the results for the presence of known terpenoid systems, and the drimane skeleton, the postulated precursor of polyathenol, was the only known skeleton which resulted. This does not prove the hypothesis of LeBoeuf, et al., but certainly helps strengthen it.

SURVEY is, however, only the barest beginning of an experiment planner, even though it has proven useful. We plan to build from this beginning toward a much more powerful system.

M. LeBoeuf, M. Hamonnière, A. Cavé, H. Gottlieb, N. Kunesch, and E. Wenkert,
Tet. Lett., 3559 (1976).

"POLYALTHENOL" $C_{23}H_{31}NO$

<u>SUPERATOMS</u>	<u>ARBITRARY NAME</u>	<u>NUMBER</u>
	IN	1
	BI	1
CH ₃ -FV	ME	1
FV-CH ₂ -FV	CH ₂	3
	CH	1

CONSTRAINTS

1) ALL FREE VALENCES BONDED TO NON-HYDROGEN ATOMS

2) GOODLIST

(EVENTUALLY	IN-CH ₂ -BI	1 TO ANY
	IN-CH ₂ -CH ₀ →0)	
	ME-(BI CH)	1 TO ANY
(EVENTUALLY	CH ₃ -CH, EXACTLY 1)	

3) GOODRINGS 2 EXACTLY 5

4) BADRINGS 3

Figure 4. Superatoms and constraints supplied to CONGEN in investigations of plausible structural alternatives to the proposed structure of Polyalthenol.















Figure 5.

Structural candidates for polyalthenol based on data given in Figure 4.

REACTION CHEMISTRY DEVELOPMENTS

1. SEPARATION FROM CONGEN - COMMUNICATION VIA FILES OF STRUCTURES.
2. ADDING CONSTRAINTS - SITE - AND TRANSFORM - SPECIFIC.
3. CONTROL STRUCTURE - RAMIFICATION
 - A. ESTABLISH RELATIONSHIPS AMONG PRODUCTS AND REACTANTS
 - B. DEAL PROPERLY WITH RANGES OF NUMBERS OF PRODUCTS
4. INTERACTION - DEVELOP MANIPULATION COMMANDS WHICH PARALLEL LABORATORY OPERATIONS, E.G., SEPARATE INTO FLASKS, TEST CONTENTS OF VARIOUS FLASKS, INCOMPLETE SEPARATIONS, ETC.
5. REPRESENTATION OF REACTIONS
6. PROSPECTIVE DETECTION OF DUPLICATE PRODUCTS BASED ON SYMMETRY PROPERTIES OF: A) STARTING MATERIAL; AND B) TRANSFORMATION.

Figure 6. Current and future direction for improvement and extension of REACT, a program for exploration of applications of reaction chemistry to structure elucidation problems.

Applications of REACT to Structure Elucidation Problems

We have recently described our initial efforts toward representation of chemical reactions and their use in structure elucidation problems [Report HPP-76-5]. These efforts provided the framework for carrying out reactions within the computer which emulate actual laboratory reactions performed on a unknown. Constraints on the numbers and identities of the products are used to constrain the reaction products and, implicitly, the starting materials. Based on the results of that work we drew up a set of steps to be carried out to provide a truly useful tool for the chemist. Although the current program can be used in applications to real problems it has some fundamental limitations which we have been working to solve. The developments we have undertaken to improve REACT are summarized in Figure 6.

We first undertook to separate REACT from CONGEN, for two reasons. One reason was due to program size. Many functions of CONGEN are not needed in REACT and become unnecessary when only REACT is being exercised. The procedures of structure generation (CONGEN) and REACT are sequential and a separate program introduces no problems. A second reason was the different uses of certain CONGEN functions in REACT. For example, the ways in which the graph matcher is used are different between the two programs, necessitating keeping two different versions around with the programs together. The separation has been accomplished. The current version of REACT is now a separate program. It communicates structural information with CONGEN via files. All interactive portions are consistent with the structural manipulation functions of CONGEN so that learning the structural language of CONGEN is sufficient to use either program.

We have also added new constraint types to the reaction to expand greatly the ways in which reactions can be defined and constrained. An example of new extensions to reaction definitions illustrates some of the new features (Figures 7-10). The reaction defined here is one which will perform a dehydration of an alcohol; the site of the reaction is defined in Fig. 7.

The transform is defined as cleavage and loss of the oxygen resulting in formation of a double bond between the two carbon atoms of the original site (Fig. 7). In this particular dehydration the chemist wished to specify a site-specific constraint. It was known that a tertiary butyl group was part of the structure, and the dehydration will be prevented if that group is in close proximity to the reaction site (i.e., in a position alpha to the carbinol carbon).

The definition of this constraint is given in Figure 8. Subsequently, this constraint ("HINDERED") is placed on BADLIST for constraints specific to the site as shown in Fig. 9. The completed definition of the reaction is summarized in Figure 10.

:EDITREACT
NAME:DEHYDRATION
(NEW REACTION)

*SITE
>CHAIN 3
>ATNAME 1 0
>HRANGE 1 1 1 3 1 3
>ADRAW

DEHYDRATION: (HRANGES NOT INDICATED)

O-C-C

>DONE

*TRANSFORM
>UNJOIN 1 2
>JOIN 2 3
>DELATS 1
>ADRAW

DEHYDRATION: (HRANGES NOT INDICATED)

C=C

>DONE

Figure 7. Definition of reaction site and chemical transform in REACT.

*DEFINE-CONSTRAINTS

:?

PLEASE ENTER ONE OF:

GRIPE BUGOUT

TRANSFORMSPECIFIC(T)

GENERAL(G)

DONE

SITESPECIFIC(S)

HALT

:SITESPECIFICNAME: HINDERED

(NEW CONSTRAINT)

(WARNING: THE FINAL CONSTRAINTS MUST HAVE AT LEAST ONE ATOM OF THE SITE)

>NDRAW

HINDERED: (HRANGES NOT INDICATED)

NON-C ATOMS: 1 0

1-2-3

>BRANCH 3 2 4 1 4 1>ADRAW

HINDERED: (HRANGES NOT INDICATED)


>DONE

Figure 8. Definition of a site-specific constraint to be applied to the reaction DEHYDRATION.

*CONSTRAINTS

:?

PLEASE ENTER ONE OF:

GRIPE

BUGOUT

ST FOR CONSTRAINTS ON STARTING MATERIAL

S FOR SITESPECIFIC CONSTRAINTS

T FOR TRANSFORMSPECIFIC CONSTRAINTS

PR FOR CONSTRAINTS ON PRODUCTS

DONE

HALT

:S

>BADLIST

BADLIST CONSTRAINTS

CONSTRAINT NAME:HINDERED

CONSTRAINT NAME:

>DONE

:DONE

Figure 9. Specification of constraint named HINDERED as a BADLIST constraint for the reaction.

*SHOW

SITE:

NAME=DEHYDRATION

ATOM# TYPE ARTYPE NEIGHBORS HRANGE

1	O	NON-AR	2	1-1
2	C	NON-AR	1 3	
3	C	NON-AR	2	1-3

DEHYDRATION: (HRANGES NOT INDICATED)

NON-C ATOMS: 1 0

1-2-3

TRANSFORM:

UNJOIN 1 2

JOIN 2 3

DELATS 1

DEHYDRATION: (HRANGES NOT INDICATED)

2=3

CONSTRAINTS:

CONSTRAINTS ON STARTING MATERIAL:

NO CONSTRAINTS

SITE-SPECIFIC CONSTRAINTS:

BADLIST CONSTRAINTS

NAME

HINDERED

TRANSFORM-SPECIFIC CONSTRAINTS:

NO CONSTRAINTS

CONSTRAINTS ON PRODUCTS:

NO CONSTRAINTS

*DONE

(DEHYDRATION DEFINED)

(DEHYDRATION ADDED TO THE REACTION LIST)

Figure 10.

Summary of the completed
definition of the
DEHYDRATION reaction.

The remaining items summarized in Figure 6 are currently under development. We are redesigning the control structure so that the scientist using the program can use intuitive concepts as commands, such as separation. To carry this out important parts of the current mechanism have to be redesigned. Although the current program can be used effectively, its non-intuitive approach to dealing with reactions yielding multiple products and subsequent separation (within the computer) and analysis of each product presents a barrier to use by a wider community. We are continuing to develop our capabilities for representing reactions to ensure that the user of REACT has a complete descriptive language with which to specify reactions. We continue to study ways to avoid duplication in carrying out reactions. We know how to implement certain of the symmetry-related constraints and will do so shortly.

CONGEN Developments

The problem solving paradigm that has emerged from DENDRAL work is the so-called "plan-generate-test" paradigm. It is based on heuristic search of a space of possible hypotheses with planning before generation of hypotheses and testing of each generated candidate.

The generator for DENDRAL, named CONGEN, is a general-purpose graph generator which produces a list of all possible graphs containing specified numbers of nodes of various types. The most important features of the generator are that the list of graphs is guaranteed to be complete and non-redundant and, equally important, that the list need not be exhaustively generated. The generator can be constrained to produce only graphs that meet specified criteria that are inferred from the initial problem data.

During the past year, CONGEN has developed along two major lines: 1) tools have been developed which will allow more efficient and "intelligent" use of substructural information supplied by the chemist; and 2) data from chemical reactions and from observed mass spectra can be used to eliminate unlikely structural candidates from a set produced by a CONGEN generation. These extensions will be discussed below.

1) Intelligent use of substructural information as constraints

There is sometimes a significant conceptual gap between the intuitive chemical phrasing of a CONGEN problem and the phrasing which is most efficient, in both computer time and storage requirements, for the program. CONGEN provides a rich language for stating structure elucidation problems in precise substructural terms. However, there are usually many ways of defining a given problem and different definitions can place widely different demands upon the program. We have a continuing interest in reducing this conceptual gap by in making CONGEN responsible for rephrasing a problem in the most efficient way, thus freeing the chemist to concentrate upon the chemical, rather than the algorithmic, aspects of a given case.

One distinction which is frequently puzzling to new CONGEN users is the one between superatoms and GOODLIST items. A superatom is a polyatomic "building block" which CONGEN joins with other superatoms and single atoms to form full


structures. GOODLIST items are substructures which are required to be present in those full structures, but they are not incorporated directly into the initial phrasing of a problem as are superatoms. Rather, their presence or absence is tested by a graph-matching routine after the structures are produced. Frequently, a great many structures produced by the structure generator are discarded by this final test and a significant amount of the program's time can be spent "shooting blanks". The concepts behind these two types of constraints - that specified substructural features must be present - are similar, but their implementations differ substantially in efficiency.

GOODLIST items cannot simply be transferred to the superatom list, though, because GOODLIST items are allowed to share atoms and bonds with other GOODLIST items or with superatoms. For example, if two substructures which are benzene rings are placed on GOODLIST, then a naphthalene derivative will be an acceptable structure even though the two occurrences of the ring have two atoms and one aromatic bond in common. Because of the building-block nature of superatoms, they may be joined to one another by additional bonds in CONGEN, but never "merged" (i.e., overlapped). Thus the price of efficiency is a more restricted interpretation of structural possibilities for superatoms.

We have developed a new procedure which captures the best of both situations. In order to incorporate a GOODLIST substructure into the problem at the earliest stage, it is necessary to find all unique ways that the given substructure can be created using parts of the existing building blocks (atoms and superatoms). This produces a set of new CONGEN problems with more or larger superatoms, each of which is easier to solve than the original one because the GOODLIST item is built-in and needs not be tested. Figure 11 shows schematically some of the ways this construction might occur: a) by bonding together two (or more) existing superatoms to create one larger one; b) by bonding additional atoms to a superatom to create a larger one; and c) by constructing a copy of the substructure from single atoms, creating a new superatom.

Figure 12 summarizes a CONGEN problem which was attempted but which could not be completed because of the unintelligent use of GOODLIST. The problem amounts to finding all ways of allocating three new bonds to the free valences (the bonds with unspecified termini) in the superatom CEMB such that the three indicated substructures are present in the final molecules. There are perhaps 10,000 unique allocations of those three new bonds, but only 7 pass the GOODLIST tests. Using GOODLIST as a post-test only, CONGEN would generate all 10,000 and discard nearly all of them, a process which would have been so lengthy that it was never completed. The constructive graph-matching routine approaches the problem in a much more efficient and chemically intuitive way: 1) there are only three places in which the first GOODLIST item can be constructed; 2) for each of these, there are four ways of constructing the second; and 3) for each of these, there are 0, 1 or 2 ways of incorporating the third. It quickly arrives at the correct set of solutions.

Most CONGEN problems contain one or more GOODLIST items which can be processed in this way, and when the constructive graph-matcher is fully integrated into CONGEN, it will make a substantial difference in its ability to use this structural information effectively.



GOODLIST:






Figure 11. Example of breaking one GOODLIST substructure into several subproblems for CONGEN, each with different superatoms.




Figure 12. Example showing the inefficiency of specifying a constraint as a GOODLIST item instead of analyzing its implications for constructing allowable chemical graphs.

2) New tools for post-pruning CONGEN structures.

From an algorithmic standpoint, CONGEN is successful if it can, in a reasonable amount of time and without exhausting storage resources, produce a list of candidate structures satisfying the chemist's constraints. However, this list is often quite large, perhaps several hundred structures, and from a chemical standpoint the problem may be far from complete. It remains for the chemist to discriminate among the candidates, eventually reducing the possibilities to just one structure. A SURVEY function is available for classifying the list into groups of chemically related structures using either pre-defined or user-defined libraries of substructural features, and this process can help the chemist perceive groups which might easily be ruled out by additional experiments. Also, the graph-matching (pruning) mechanism of CONGEN allows him to express, in terms of substructural tests on the candidates, new data which he gathers on the unknown. These are both important aids in dealing with a list of candidates, but are restricted to tests which can easily be phrased purely in terms of structural features of the candidates themselves.

There are two informative sources of data which cannot always be phrased in this way: 1) structural features observed in products of the unknown when it undergoes simple chemical reactions; and 2) empirical spectroscopic measurements on the unknown which cannot be interpreted unambiguously in precise structural terms. During the past year, we have made progress in utilizing such information. The program REACT addresses the first problem while MSRANK concerns the second, in the context of mass spectrometric observations.

2.1 REACT

This program [see Report HPP-76-5] has two basic goals: 1) to provide the chemist with a computerized language for defining graph transformations and applying them to structures, thus simulating chemical reactions; and 2) to automatically keep track of the interrelationships between structures in a complex sequence of reactions so that whenever structural claims are made ruling out structures at one level, the implications in terms of structures at other levels can be traced. During the last year some progress has been made toward both of these goals.

EDITREACT, the reaction-editing language, has been extended to allow the user to define subgraph constraints which apply relative to a potential reaction site rather than to the molecule as a whole. For example, in the present version of REACT, we can say either that a hydroxyl group (OH), if present anywhere in the reactant molecule, would inhibit the reaction, or that such inhibition would take place only if the OH group is adjacent to the reaction site. Such site-specific constraints, applied either before or after the transformation (i.e., reaction) has been carried out on the site, are critical to the detailed description of real chemical reactions. The inclusion of this facility in REACT substantially increases its usefulness in real-world chemical problems.

The bookkeeping problem has undergone a complete reconceptualization in the past year, the purpose being to mimic more closely the actual steps taken by a chemist in the laboratory. In the initial implementation, a set of products arising from the application of a given reaction to a given starting structure

could be subjected to a multi-level classification which grouped the products based upon user-defined substructural constraints. Each of these classes had an associated minimum and maximum number, representing the numbers of products which were allowed to be members of the class. Any starting materials whose products could not satisfy these conditions were removed from the list of candidates. Structures in any class could be further reacted, their products classified, and so on. This treatment of bookkeeping was sufficient for stating many chemical problems. For example, suppose a chemist knew that a particular reaction on an unknown compound yielded two carbonyl compounds (i.e., containing $C=O$), at least one of which was an ester ($-O-C=O$). He could define a product class CARBONYL using the $C=O$ substructure with a minimum and maximum of two products. He could then define a sub-class of CARBONYL called ESTERS using the substructure $-O-C=O$ with a minimum of one and a maximum of two products. The program would automatically use this information to eliminate candidate starting structures which could not give the indicated product distribution with the given reaction.

There are chemical problems, though, for which the above scheme is too rigid. For example, suppose a reaction gives several products, two of which are isolated and labelled P1 and P2. Suppose that only a small amount of P1 is available so only mass spectroscopic measurements are practical. Suppose also that a deuterium-exchange experiment shows that P1 has two exchangeable protons (say, either N-H or O-H). P2 shows a strong carbonyl absorption in the IR. P1 might also contain a carbonyl group, but that was never determined, and neither was the number of exchangeable protons in P2, which could be two. No matter how one attempts to use the above-described classification system, one cannot express this information accurately.

In the new approach, for which the algorithmic design has been completed, one is allowed to express data in a much more natural sequence which parallels the experimental steps. The first experimental step after a reaction is usually the separation and purification of products. An analogous step is to be included in REACT, in which the separation amounts to the setting up of a specified number of labelled "flasks" (analogous to the labels P1 and P2 in the above example) each of which is ultimately to contain a specified number (usually 1) of the products. As experimental data are gathered on each real product, corresponding substructure constraints are attached to the corresponding flask in the program. As each such assertion is made, the bookkeeping mechanism verifies that, for a set of reaction products from a given starting material, there is at least one way of distributing them among the flasks such that each product satisfies the constraints for its flask. If this test is ever violated, the starting material is removed as a candidate structure. Flasks containing more than one product may be further separated into "subflasks" to any level, and the contents of any flask may be made to undergo further reactions. This capability, the reacting of flask contents, is analogous to common laboratory procedures in which incomplete separations of products are encountered. Dealing with such situations adds considerable complexity to the bookkeeping mechanism, because the contents of a flask may be ambiguous to the program when the reaction is applied. REACT must keep track of all possible structures which might, based on the current flask constraints, occupy the reacting flask. If such a reaction fails (because the products did not satisfy the constraints specified for them), REACT does not eliminate the starting structure entirely, but notes that the structure may not occupy that flask in future flask-allocation tests.

2.2 MSRANK

This program is an outgrowth of MSPRUNE described in last year's annual report. It is a combination of a predictor which uses a very simple theory of mass spectrometry to predict the spectra of candidate structures, and an evaluation function which compares the predictions with the observed spectrum of the unknown, assigning a goodness-of-fit score to each candidate. The candidates are then sorted based upon how well they match the observations. The basic concept here is not a new one to the DENDRAL project [see, for example, Buchanan, et al. in Machine intelligence 4 (Meltzer & Michie, eds., Edinburgh Univ. Press, 1969)], but there are some new aspects to the problem when viewed in the overall CONGEN context.

Because of the wide variety of structural types which can be produced by CONGEN, it is necessary for MSRANK to use a very general model of mass spectrometry. The best predictive theories of mass spectrometry are limited to families of closely related structures (i.e., class specific theories), and the Meta-DENDRAL program is designed to help in discovering such theories. There are very few general principles upon which to draw in predicting mass spectra, though, so MSRANK is limited to only the most approximate kinds of evaluation functions. One principle which we noticed being used by practicing mass spectrometrists was: of two candidate structures for an unknown, the most likely structure is the one which explains the observations most "simply" - i.e., with the fewest complex explanations involving many bond cleavages and the transfer of many hydrogen atoms. The evaluation function used by MSRANK is based on a quantitation of this principle.

MSRANK is quite new and we have not yet had sufficient experience with it to evaluate its overall usefulness. By using only unit plausibilities for selected characteristics of the mass-spectral cleavages, we are able to duplicate earlier results obtained with the predictor/comparator functions applied to mono- and di-ketoandrostanes. These tests serve to check the accuracy of the MSRANK program. We are now doing a systematic study of various classes of compounds by ranking the spectrum of a known structure against a CONGEN-generated list of structures which contains the correct one among several which are closely related.

Stereochemistry in CONGEN

We have started the complex task of giving CONGEN the capability of recognizing stereochemical features of molecules and using stereochemical information in structure determination. The ability to recognize stereochemical features would allow, for example, the generation of all stereoisomers of a given topological structure with or without constraints. The ability to use stereochemical information would allow the determination of constraints on stereoisomer (and topological isomer) generation caused by, for example, partial knowledge of relative or absolute stereochemistry of structural fragments, knowledge of overall molecular chirality (or lack of), absolute and relative

stereochemistry from circular dichroism measurements, and so forth. Thus far, only the topological information (constitution) has been recognized and used by CONGEN.

The first stage of this development is to produce a program which generates all the stereoisomers of a given topological structure. This program will be placed at the end of the existing CONGEN program. The present report describes the development of the theory and algorithm for stereoisomer generation and the progress on the programming of this algorithm.

The GC/HRMS DATA SYSTEM

New Developments

In addition to upgrading old versions of the high resolution system, work is being done on creating a low resolution system for the MAT 711. The ultimate aim is collect data that can be run through CLEANUP, a program that resolves multiple spectra under a single GC peak, and cleans up the final spectra. The problem with the current system is that we cannot scan fast enough to provide CLEANUP the data it needs. The high resolution system requires resolution good enough to separate sample peaks from the reference peaks. If the scan is sped up past a certain point, SAMRUN can no longer separate the peaks, and therefore cannot calibrate the run. At the same time, CLEANUP requires at least 7 spectra across a GC peak be taken to insure resolution of multiple spectra. The fundamental problem then is that an alternate method of calibrating the mass spectrum, without using known calibration peaks, must be found before scan speeds required by CLEANUP can be achieved. The most direct solution to this is to directly measure the magnetic field strength of the instrument, and using it to calculate the mass that is being observed. To do this we inserted a hall probe between the poles of the magnet, and connected it to the data acquisition system on the PDP-11/20.

The main problems with the hall probe are as follows: 1) to make sure that the ion reading and the hall probe reading are simultaneous 2) to insure that the correct hall reading can be assigned to the correct ion reading 3) to determine the reproducibility of hall readings versus mass being observed in both dynamic (scanning) and static situations and 4) to decide if the probe has the speed and accuracy to calibrate the instrument. The first two problems are a matter of hardware. The configuration of the original data collection system is as follows: the ion detector goes to an A/D converter, which is connected to a DMA. The DMA is on an 11/20, which has a data collection system, SAQMON, running. This performs various low level filtering and buffering operations. The DMA is actually a low level processor which counts the number of samples taken, stores them into successive memory locations, and interrupts the central processor when a block of data has been collected. The timing of the sample collection is controlled by a quartz crystal clock. On each timing pulse, a signal is sent to the A/D on the ion detector to convert that value to a digital number. To

accommodate the hall probe, the DMA was modified so that on the timing pulse, the start signal is sent simultaneously to both the A/D on the ion detector and the A/D on the hall probe. The DMA then services both of the A/D's, and stores the readings in successive memory locations. The net result is that when the DMA interrupts the central processor, the block of data is a set of pairs of readings, an ion reading and the hall reading for that time. This solves both of the first two problems, since we now have the ion reading and the hall reading connected both in time and location.

The second two problems, testing the reliability and reproducibility of the hall probe, requires new software. We are currently modifying portions of the calibration mechanism of the high resolution system to calculate masses for a large number of hall readings.

META DENDRAL

The success of any reasoning program is strongly dependent on the amount of domain-specific knowledge it contains. This is now almost universally accepted within AI, partly because of DENDRAL's success. Because of the difficulty of extracting specific knowledge from experts to put into the program, many years ago we began to explore the problems of efficiently transferring knowledge into a program. We have looked at two alternatives to "hand-crafting" each new knowledge base: interactive knowledge transfer programs and automatic theory formation programs. In this enterprise the separation of domain-specific knowledge from the computer programs themselves has been a critical component of our success.

One of the stumbling blocks with the interactive knowledge transfer programs is that for some domains there are no experts with enough specific knowledge to make a high performance problem solving program. We were looking for ways to avoid forcing an expert to focus on original data in order to codify the rules explaining those data because that is such a time-consuming process. Therefore we began working on an automatic rule formation program (called Meta-DENDRAL) that examines the original data itself in order to discover the inference rules for that part of the domain.

The problem solving paradigm for Meta-DENDRAL is also the plan-generate-test paradigm used in Heuristic DENDRAL. In this case one part of the program (RULEGEN) generates plausible rules within syntactic and semantic constraints and within desired limits of evidential support. The model used to guide the generation of rules is particularly important since the space of rules is enormous. The planning part of the program (INTSUM) collects and summarizes the evidential support. The testing part (RULEMOD) looks for counterexamples to rules and makes modifications to the rules in order to increase their generality and simplicity and to decrease the total number of rules.

Meta-DENDRAL successfully formulated rules of mass spectrometry that were new to the science. These rules, along with a discussion of the methodology,

were published in the scientific literature [Report HPP-76-4]. The program was tested to see if it could rediscover the rules of mass spectrometry for two classes of chemical compounds that were already well understood (amines and estrogenic steroids). Then it was applied to three classes of compounds whose mass spectrometry was not as well known (mono-, di-, and tri-ketoandrostanes). The program produced three sets of rules that explained much of the significant data for these classes. The time for manual rule formation for these data was estimated to be several months.

Progress was made on generalizing the Meta-DENDRAL program, and rules for a new domain were successfully discovered by the program. A scientific paper on this application was submitted for publication [Report HPP-77-4]. The new application was learning rules for interpreting signals from C13-NMR spectroscopy. The instrument produces data points in a bar graph in response to the resonance of each carbon-13 nucleus in the sample. The rules describe an environment of a C13 atom and predict a resonating frequency range for every atom that matches the description. The Meta-DENDRAL program needed some modification because the rules are predicting ranges of data points, and not precise processes, as for the mass spectrometry version.

The RULEGEN component of Meta-DENDRAL was demonstrated to work with its heuristic search paradigm. Guidance from a model of mass spectrometry is an important feature of RULEGEN. Also, the program uses problem data for pruning possible rules (and all more specific rules formed from those). The amount of data examined during the search is very large and the space of rules is immense, so the search needs to be rather coarse in order to produce plausible, but not necessarily optimal, rules.

The RULEMOD program for "fine-tuning" Meta-DENDRAL's newly-discovered rules was finished. This program provides a number of important subtasks, including merging similar rules, making rules more specific or more general, and filtering out the weakest rules. RULEMOD checks for counterexamples to rules and uses this information in all of the named tasks. Because of the expense of computing counterexamples to possible rules, this computation is delayed until Meta-DENDRAL has a set of plausible rules, rather than computing counterexamples on each possible rule examined in the search of the rule space.

A report was written on the AI methodology underlying Meta-DENDRAL. The major idea developed in this report is that knowledge of the domain can be used effectively to guide a learning program. The major difference between Meta-DENDRAL and statistical learning programs is that Meta-DENDRAL uses a strong model of mass spectrometry, including any assumptions the user cares to make about the domain, to guide the formation of explanatory rules.

C13 NMR SPECTROMETRY

¹³C NMR was selected as a new application area for the rule formation program, Meta-DENDRAL. The algorithms used for mass spectrometry rule formation

were extended to ^{13}C NMR and used to obtain a set of rules for These two classes and acyclic amines. These two classes were chosen since compounds in these classes are known to show a strong correlation between structural environment and shift. Thus, the programs could be tested knowing that the underlying basis for the form of the rule was valid.

The form of the rule is

substructure ---> shift range.

A sample rule generated is

$\text{C}-\text{C}^*-\text{C}-\text{X}-$ ---> $19.85 \leq (\delta_{\text{C}}) \leq 21.3$.

The asterisk in the substructure description denotes the atom for which the shift is predicted. Only topological descriptors were used to construct the substructures. The addition of stereochemical terms is a topic of current work.

It was necessary to change RULEGEN so that the left-hand sides of rules were expanded outward from a carbon atom rather than from a bond. The right-hand side of the rule is associated with a range rather than a precise mass as in the mass spectrometry program. This modification also required changes in the rule search procedure. The user sets two parameters which guide the rule search. These parameters are MINIMUM-EXAMPLES which requires each rule to explain a given number of peaks in the training set and MAXIMUM-RANGE which defines the acceptable shift range for a rule. These parameters regulate the degree of specificity or generality of the rules.

From the set of rules generated a subset is selected corresponding to the "best" set which still covers all the training set data. The best rule is selected by calculating

$(\text{number of peaks predicted} / (\text{range} ** 2))$.

Data which are predicted by the best rule are removed and the next best rule is found for the remaining data using the criterion given above. This process is repeated until all data are explained.

In order to test the informational content of the rules generated a second program was written which applied the rules to a list of candidate molecules and ranked the molecules. Firsts, all possible structural isomers for a given empirical formula were generated using CONGEN. The rules were applied to each of the possible isomers and spectra were predicted. The predicted spectra were compared to that of a known spectrum from a compound with the same empirical formula. The structural isomers were ranked according a comparison score to determine how well the correct compound was distinguished from its isomers, on the basis of the predictive rules.

The details of the generation of rules and the use of rules for structure selection can be found in a paper recently submitted for publication [Report HPP-77-4]

The ^{13}C NMR rule formation program was applied to a set of paraffins and acyclic amines. The program generated 138 rules to cover 435 data peaks. The rules generated were applied in a structure selection test for the structural isomers of C_9H_{20} and $\text{C}_6\text{H}_{15}\text{N}$. No structures with these empirical formulas were included in the training set. Twenty-four C_9H_{20} and eleven $\text{C}_6\text{H}_{15}\text{N}$ ^{13}C NMR spectra were available to act as unknowns in the structure selection test. The results of the structure ranking applied to these spectra are shown below.

EMPIRICAL FORMULA	NUMBER OF CANDIDATE ISOMERS	NUMBER OF CANDIDATES RANKING			
		1st	2nd.....6th.....9th		
C_9H_{20}	35	20/24	3/24		1/24
$\text{C}_6\text{H}_{15}\text{N}$	39	8/11	2/11	1/11	

The performance of the rules in discriminating among similar structures not included in the training set data demonstrated the content of the rules.

RECENT PUBLICATIONS

(Only publications related to computers in chemistry are shown.)

- HPP-76-1 D.H. Smith, J.P. Konopelski and C. Djerassi, "Applications of Artificial Intelligence for Chemical Inference. XIX. Computer Generation of Ion Structures", *Organic Mass Spectrometry*, 11: 86, (1976).
- HPP-76-2 Raymond E. Carhart and Dennis H. Smith, "Applications of Artificial Intelligence for Chemical Inference XX. Intelligent Use of Constraints in Computer-Assisted Structure Elucidation", *Computers In Chemistry* (in press).
- HPP-76-3 C.J. Cheer, D.H. Smith, C. Djerassi B. Tursch, J.C. Braekman and D.

Daloze, "Applications of Artificial Intelligence for Chemical Inference XXI. Chemical Studies of Marine Interbrates - XVII. The Computer-Assisted Identification of [+-]Palustrol in the Marine Organism *Cespitularia* sp., aff. *subviridis*". *Tetrahedron*. 32:1807, Pergamon Press, (1976).

- HPP-76-4 B.G. Buchanan, D.H. Smith, W.C. White, R.J. Gritter, E.A. Feigenbaum, J. Lederberg, and Carl Djerassi, "Application of Artificial Intelligence for Chemical Inference XXII. Automatic Rule Formation in Mass Spectrometry by Means of the Meta-DENDRAL Program", *Journal of the American Chemical Society*, 98: 6168 (1976).
- HPP-76-5 T.H. Varkony, R.E. Carhart and D.H. Smith, "Applications of Artificial Intelligence for Chemical Inference XXIII. Computer-Assisted Structure Elucidation. Modelling Chemical Reaction Sequences Used in Molecular Structure Problems", in "Computer-Assisted Organic Synthesis", W.T. Wipke, Ed., American Chemical Society, Washington, D.C., in press.
- HPP-76-6 D.H. Smith and R.E. Carhart "Applications of Artificial Intelligence for Chemical Inference XXIV. Structural Isomerism of Mono and Sesquiterpenoid Skeletons 1,2-", *Tetrahedron*, 32:2513, Pergamon Press (May 1976).
- HPP-76-10 Bruce G. Buchanan and Dennis Smith, "Computer Assisted Chemical Reasoning", in *Proceedings of the III International Conference on Computers in Chemical Research, Education and Technology*, Plenum Publishing, (1976).
- HPP-77-4 T.M. Mitchell and G.M. Schwenzer, "Applications of Artificial Intelligence for Chemical Inference. XXV. A Computer Program For Automated Empirical ¹³C NMR Rule Formation", (Submitted to JACS, January 1977).
- HPP-77-6 Bruce G. Buchanan and Tom Mitchell. "Model-Directed Learning of Production Rules", Submitted to the *Proceedings for the Workshop on Pattern-Directed Inference Systems in Hawaii*, (February, 1977). (STAN-CS-77-597)
- HPP-77-11 Dennis H. Smith and Raymond E. Carhart, "Structure Elucidation Based on Computer Analysis of High and Low Resolution Mass Spectral Data". *Proceedings of the Symposium on Chemical Applications of High Performance Spectrometry*. University of Nebraska, Lincoln, (in press).

II. INTERACTION WITH THE SUMEX-AIM RESOURCE

The number of persons experimenting with CONGEN has grown as a result of both the continuing practice of issuing an "invitation for program trial use" at the conclusion of publications, as well as continuing personal contact between

Dendral project members and potential program users. Three categories of users make up this group:

Chemists Using Exported Programs

The part of CONGEN responsible for teletype output of chemical structures (the DRAW program) is coded in Fortran. Since the paper describing this program appeared in print [R. Carhart, JACS, 16:82, 1976], we have exported the program to half a dozen sites, ranging from Japan, across North America, to England. Similarly, the entire CONGEN program, is largely coded in Interlisp and SAIL, and has been exported to a collaborator in England who is very interested in the methods and programming techniques employed in coding the program. Another program which we have exported for use by other chemists is the PDP-11 CLEANUP program which was described in ANALYTICAL CHEMISTRY [48:1368, 1976]. This program "cleans up" new GC/MS data to eliminate noise peaks and to separate the data associated with components in the mixture.

In each case, the requestors were provided with an initial choice of format options from which they could select the one most suitable for their computer installation. They were asked to send a 2400 foot reel of magnetic tape appropriate to the selected format option. The programs were written on the tape and returned to them along with a brief written explanation of program organization. Accurate records are kept of who has received the programs, so that omissions and errors can be corrected by mail at a later date, if ever necessary.

1. Dr. James F. Elder, Dow Chemical U.S.A., Midland, Michigan.
2. Dr. Robert M. Supnik, Massachusetts Computer Associates, Inc., Wakefield, Massachusetts.
3. Mr. Dan Pearce, Orange County Sheriff-Coroner Department, Santa Ana, California 92702
4. Dr. H. J. Stoklosa, Central Research & Development Department, E. I. du Pont de Nemours & Company, Wilmington, Delaware.
5. Dr. Douglas W. Kuehl, Environmental Research Laboratory-Duluth, Duluth, Minnesota.
6. Dr. Richard A. Graham, Food Sciences Laboratory, U. S. Army Natick Laboratories, Natick, Massachusetts.
7. Dr. Walter M. Shackelford, United States Environmental Protection Agency, Environmental Research Laboratory, Athens, Georgia.
8. Dr. Richard Gans, Chemical Research Division, American Cyanamid Company, Bound Brook, New Jersey.
9. Dr. John C. Marshall, Department of Chemistry, the University of North Carolina, Chapel Hill, North Carolina.
10. Dr. Graham S. King, Department of Chemical Pathology, Queen Charlotte's Hospital for Women, London, England.

11. Dr. J. Wyatt, Chemistry Division, Naval Research Laboratory, Washington, D. C..
12. Dr. Gareth Templeman, Research and Development Laboratories, The Pillsbury Company, Minneapolis, Minnesota.
13. Dr. J. B. Justice, Department of Chemistry, Emory University, Atlanta, Georgia.
14. Dr. Thomas Knudsen, Northrop Services, Environmental Sciences Group, Research Triangle Park, North Carolina.
15. Dr. Ingolf Meineke, Fachbereich Chemie, Philipps Universitaet, Lahnberge, West Germany.
16. Dr. M.A. Shaw, Unilever Research, Port Sunlight Laboratory, Wirral, Merseyside, England.
17. Dr. Ernst Weber, Varian MAT, Bremen, West Germany.
18. Paul V. Fennessey, Department of Pediatrics, University of Colorado Medical Center, Denver, Colorado.
19. R. G. A. R. Maclagan, Department of Chemistry, University of Canterbury, Christchurch, New Zealand.
20. James E. Oberholtzer, Arthur D. Little, Inc., Cambridge, Massachusetts.
21. F. Street, AEI Scientific Apparatus Limited, Manchester, England.

Remote Users of SUMEX

Due to the fact that the SUMEX computer is available via both the TYMNET and ARPANET communication networks, it is possible for scientists in many parts of the world to directly access the Dendral programs on SUMEX. Primary usage is centered on CONGEN, although INTSUM is beginning also to gain a following. Although access points to SUMEX are widespread, they frequently are not diverse enough to accommodate the dispersed group of scientists who have expressed an interest in using one of the Dendral programs. For example, Dr. Joseph Baker of the Roche Institute of Marine Pharmacology in Dee Why, Australia, is looking at the possibility of accessing SUMEX by using International Direct Distance Dialing (IDDD).

Chemists Communicating by Mail

Many Scientists interested in using DENDRAL programs in their own work are not located near a network access point. Users of this type choose to use the mail to send details of their structure elucidation problem to a Dendral Project collaborator at Stanford.

Chemical Problems Posed to CONGEN

Following is a list of CONGEN users, and a brief summary of their program interests during the past year.

1. Dr. Roger Hahn, Syracuse University. While at Stanford he used CONGEN to help solve the structures of photoproducts by obtaining all possibilities under available constraints and designing NMR experiments to differentiate the possibilities. This work will be published soon.
2. Dr. William Epstein, University of Utah. During a demonstration of CONGEN, he posed a problem to verify that the structural possibilities he determined for an unknown were in fact all possibilities. The structure of methyl santolate has been published (see Epstein, et al., J.C.S. Chem. Commun., 590 (1975)).
3. Dr. Clair Cheer, University of Rhode Island. While on sabbatical at Stanford, Dr. Cheer has worked on a number of structure elucidation problems using CONGEN including Briareine D and [+-]-Palustrol (Cheer et al., Tetrahedron Letters, 1807 (1976)). Work is continuing on the structure of another marine natural product, presumably a cembrenolide, for which there are currently seven possibilities.
4. Dr. Jerrold Karliner, Ciba-Geigy Corporation. Dr. Karliner has solved several structural problems using CONGEN, including material with flame retardant properties, an impurity in a production sample and nitrogen heterocycles being investigated for pharmacological activity. CONGEN enabled reduction of the number of possibilities to the point where subsequent experiments led to unambiguous structural assignment.
5. Dr. Gino Marco, Ciba-Geigy Corporation. He has used CONGEN to help solve structures of conjugates of pesticides with sugars and amino acids.
6. Dr. Milton Levenberg, Abbott Laboratories. He has worked on the structure of a compound with mild antibiotic activity, isolated from a fermentation broth. There are currently ten structural possibilities, reduced to that number from the 33 initially determined using CONGEN by additional experimental data.
7. Dr. David Pensak, DuPont. He is currently learning to use CONGEN and plans to evaluate its utility for structural problems of some of his coworkers.
8. Dr. Douglas Dorman, Eli-Lilly. He is using CONGEN to assist in structure elucidation of metabolites of microorganisms shown to have pharmacological activity. He has worked on five such problems, including a current one where the developing MSPRUNE capabilities are being used.
9. Dr. L. Minale, Napoli, Italy. We have worked with him by sending him

- structural alternatives for proposed structures for some marine natural products (Pallescensins, Tetrahedron Letters, 1417 (1975)) and cyclic diethers from the lipid fraction of a thermophilic bacterium (J. C. S. Chem. Commun., 543 (1974)).
10. Dr. K. Nakanishi, Columbia University. We have worked with him by sending him structural possibilities for termite defense compounds (structure finally solved by X-ray crystallography). This trial plus a live demonstration to one of his students has resulted in efforts toward continued collaboration on other insect defense secretions and exploration of the possibility of his direct access to SUMEX.
 11. Dr. L. Dunham, Zoecon Corporation. We have collaborated with him on the use of INTSUM for mass spectral fragmentation studies of insect juvenile hormones.
 12. Dr. A. G. Gonzales, Tenerife, Spain. We have recently sent him structural alternatives for constituents of Laurencia Perforata (Tetrahedron Letters, 2499 (1975)), and expect to continue discussions on the structures of these compounds.
 13. Dr. T. Irie, Sapporo Japan. We have recently sent him structural alternatives to published structures on constituents of Laurencia Glandulifera (Tetrahedron Letters, 821 (1974)) and expect to continue discussions on this problem.
 14. Dr. C. J. Persoons, Delft. We have corresponded with him on structural alternatives for cockroach sex pheromones (Periplanone-B (Tetrahedron Letters, 2055 (1976))), and he has agreed to further collaboration on new problems.
 15. Dr. F. Schmitz, University of Oklahoma. We explored for him structural alternatives for an unknown diterpenoid hydrocarbon. We obtained 25 possibilities, of which only four obeyed the isoprene rule.
 16. Dr. J. Baker, Roche Institute of Marine Pharmacology, Australia. We plan collaboration with Dr. Baker on the sterol fractions of various marine organisms and are exploring ways for him to access CONGEN.
 17. Dr. E. VanTamelen, Stanford University. We have used the developing reaction features of CONGEN to explore structural possibilities for both chemical and biogenetic cyclization products of squalene-oxide congeners. We have suggested alternatives to proposed structures and helped to design experiments to differentiate them.
 18. Dr. J. C. Braekman, Brussels. Dr. Braekman visited Stanford as a part of continuing collaboration in marine chemistry with Dr. Tursch's group. While at Stanford he explored use of CONGEN for use in current problems in marine natural products, and worked on the problems of Drs. Irie and Gonzales (see above). He is currently exploring access to CONGEN from Brussels, via TYMNET.

Use of CONGEN by working scientists has turned up one major area in which additional information to the user was thought to be necessary. CONGEN users unanimously indicated their desire for a method of determining what percentage of the whole problem was solved at any moment, i.e., total number of possible structures is represented by the number already generated. In a prototype system we have implemented the Cntrl-I and Cntrl-S user information interrupts, to show how far CONGEN has progressed. If, for example, someone who has generated 357 structures is told that this indicates that they have generated 1 percent of the total possible structures, they immediately know that they do not want to finish generating all the structures. Even if there were enough space, 40,000 structures would be far more than they would want to see.

We implemented another user-oriented facility for an invited paper presented at the 172nd American Chemical Society meeting, in August of 1976. Special features were added for a character-oriented, screen-addressable CRT terminals to give users an informative visual interface to CONGEN, an otherwise complex. The dynamic field of view provided by this type of terminal was used to advantage to give the chemist-user a continuous, graphic summary of both the information he has supplied to the program and the dynamic use of that information by the program.

INTERACTION WITH OTHER SUMEX-AIM PROJECTS

We have had numerous discussions with Prof. Todd Wipke's research group in meetings of our combined groups. Because the problems of manipulating chemical graphs are much the same for both groups, frequent discussions are mutually advantageous.

Almost daily contact with other Stanford-based projects provides new ideas and programming assistance. In particular, there is considerable interaction with members of the MYCIN, MOLGEN and Protein Crystallography projects. Many of our experiment planning ideas have come from discussions with the MOLGEN group. Our ideas about explaining a program's reasoning are derived from the success of MYCIN's explanation package. And our ideas about integrating multiple sources of knowledge in data interpretation have been enhanced through discussions with the Protein Crystallography group. The large number of excellent INTERLISP programmers in all these groups provides a pool of programming expertise that we draw on frequently also.

We are collaborating with Dr. Robert Lindsay on a monograph about the DENDRAL programs, with most of our interaction and all our text preparation taking place over the SUMEX system. We have also discussed helping Dr. Lindsay with a knowledge-based reasoning program to help pathologists at the University of Michigan.

6.1.2 HYDROID PROJECT

HYDROID - Studies in Distributed Processing and Problem Solving

Prof. Gio Wiederhold
Computer Science and Electrical Engineering
Stanford University

I. Summary of Research Program

A. Technical Goals

The objective of this research is the development of a methodology for the analysis and implementation of alternatives in distributed processing and problem solving. One of the primary reasons for interest in this area is its potential to break through the speed limitation barriers imposed by uniprocessing systems. If such a breakthrough can be achieved then the viability of the methods being developed by other projects using the SUMEX-AIM resource will be enhanced.

The rapid development of microprocessor and communications technology has given rise to a large number of proposed implementations of networks employing multiple processors. The computations to which these distributed systems are to be applied include heuristic decision-making problems, mathematical modelling, data reduction, and database search, as well as general purpose multi-access computing. There is however a lack of an adequate global understanding of the computational tradeoffs implied by network architectures.

In order to complement the experimental results of other investigators and broaden their applicability to the system-design decision-making process, we are developing a general framework for the study of processor interaction in distributed processing systems. The framework consists of rules to obtain parameters from programs which specify the computations, rules to parameterize descriptions of networks of processors, and procedures to calculate expected system performance from these parameter sets. The framework is to be sufficiently powerful so that, when it is validated, the methods will be able to assist in the a priori assessment of the potential performance of new system alternatives or of systems with improved system components.

One of the primary tools we are using to analyze the interaction between computations and distributed processor networks is simulation. The behaviour of processor network nodes, interprocessor control and task flow, and problem decomposition all require simulation at different levels of abstraction. Analytic queuing models may provide insight into relationships in networks, but are not adequate to provide quantitative results. Simulation is not seen as the end product of the study, but as a means to develop and assess the validity of our model of the interaction of computations and processor network architecture. Where possible, mathematical results will be used to assess the validity of model simulations.

A number of large computational applications are being analyzed in order to assess their potential for decomposition into modules for distributed processing. The current candidate applications are:

- a) Programs which use heuristic methods in decision-making. Heuristic programs frequently employ recursive decomposition of problems into subsidiary problems which themselves may be suitable for distributed processing.
- b) Programs which use multi-faceted databases to retrieve and abstract information. The process of intelligent data retrieval and analysis often depends on data or knowledge sources which are being maintained at geographically distributed processing sites.
- c) Programs which acquire data from multiple, possibly dissimilar, sensors and attempt to reduce this data to simpler hypotheses.
- d) Programs which solve large numerical problems, such as those found in image processing applications.

Parameters which describe the computations to be simulated include:

- a) The computational kernel size: the cycle and memory demand of a computational unit between interprocessor reference requirements.
- b) The computation definition message size: the amount of data required to transmit sufficient information to initiate a computational kernel.
- c) The database size: the amount of data or program text required to sustain a computational kernel, and its availability and residence in the network.

The behaviour of the system can be varied through the adjustment of other parameters. These parameters may be set to reflect the architecture of specific hardware systems, or may be varied to obtain optimum performance. In addition to obvious parameters (as the number and power of the processors), we expect the following parameter types to be important in developing an understanding of the spectrum of distributed processor architectures:

- a) Interconnection density. As the density decreases, the message delay and congestion increase. This parameter will provide a high level abstraction of multi-processor connectivity schemes. Geographical distribution will increase message delay and transmission cost.
- b) Computational locality. A high degree of locality (of database or procedural information in the network) will enhance the probability that relevant knowledge exists in closely linked nodes, thus counteracting the effects of a low interconnection density.
- c) Database viscosity. A database, including the programs required to carry out the computations at a node, may be more or less fixed to one specific node. This therefore encourages the use of certain nodes for specific functions. Many current processor networks are completely rigid in this sense, and for these networks optimal initial program and database

allocations may be determined. However, we hypothesize that a greater degree of dynamic resource allocation is desirable to cope with changing loads and in order to enhance reliability. For this reason this parameter needs to be included.

- d) Redundancy. In order to assess the cost and benefits in terms of responsiveness and reliability, the redundancy of database and computations will also be made a parameter. In order to utilize the redundancy well, the computational resources (programs or data) which effect system performance most must be identifiable.
- e) Error rate. In order to test the effectiveness of reliability strategies, node and communications channel failures will be simulated.

An important aspect of this model is that we intend to keep the abstractions at a sufficiently high level to allow analytic and intuitive verification of the model behaviour when applied to well understood computations. Computations have been mapped into specific parallel machines, but these results are not easily transferred to new architectures. The distributed processor systems now being built may have characteristics with unpredicted effects on system behaviour. We expect to be able to use the model to find potential bottlenecks, which then will define areas where extra design attention has a high payoff.

We do not intend to build hardware which is based literally on the abstract model. We hope to verify results obtained from the model using existing distributed processor systems and, assuming that our model (with appropriate parameters describing the load and architecture) matches the given system, be able to advise on system utilization or development aspects. A local resource of this type may be the Stanford I processor, now being built under ERDA sponsorship. In addition, if we determine that a certain, yet untried, architecture is promising, we would like to encourage and participate in its implementation.

B. Medical Relevance and Collaboration

Many applications at SUMEX consume large quantities of computational resources. The use of multiple distributed processors may provide a means to gain the required processing capabilities in an economic manner. In this sense the medical relevance of this study is indirect. We are attempting to develop tools which will be of use in medical computation problems.

Our studies in distributed data base applications have a more direct medical relevance. To this end, we are maintaining contact with Dr. Jim Fries, whose ARAMIS database network collects data for the analysis of disease progress and treatment efficacy in rheumatoid arthritis from a variety of institutions. Sharing of data to provide a broader base for analysis is also a feature of programs in cardiology and oncology in which physicians at Stanford participate. In each of these instances the distributed nature of the data resources leads to differences in the meaning of data items, so that simple aggregation of the data may not be valid. Distributed processing may provide a powerful alternative.

C. Progress Summary

The HYDROID project got underway in the fall of 1976. We have been involved since that time in developing a basic understanding of important problem areas in distributed processing and problem solving.

A weekly research seminar, begun in Dec. 1976 has brought together members of the faculty and students from a variety of disciplines, and has included several speakers from application areas where distributed processing may be beneficial.

We have developed a formalism in which to express the control of distributed problem solving in loosely-coupled processor networks. This CONTRACT NET protocol makes the cost of interprocessor interactions explicit. It is this cost which appears to generate one of the performance boundaries for distributed processor systems.

We have written a basic simulator with which to investigate the merits of the formalism together with problem solving methods applicable in the distributed processing environment. To this end the simulator is currently being tested with small search problems as a means of determining the necessary information that must be transferred from node to node in a distributed processor system for such problems together with the advantages to be accrued via a distributed approach. The simulator is being developed to cover a greater variety of computational interactions.

D. Publications

- 1) H. Garcia-Molina and Gio Wiederhold, "Application of the Contract Net Protocol to Distributed Data Bases", HPP-77-21, Heuristic Programming Project, Stanford University, April 1977.
- 2) R. G. Smith, "The Contract Net: A Formalism for the Control of Distributed Problem Solving", HPP-77-12, Heuristic Programming Project, Stanford University, February 1977 (also submitted to the Fifth International Joint Conference on Artificial Intelligence).

II. Interactions with SUMEX-AIM

SUMEX-AIM currently provides all computing resources for the project. We thus enjoy a high degree of interaction with other projects involved in the problems which result from construction of large programs. Other points of contact are related to the use of the same programming languages as well as the abundance of AI expertise residing around the resource. This latter point is

especially important considering that one of our aims is discovery of suitable mappings of well understood AI methods onto highly parallel asynchronous processor networks.

SUMEX-AIM is also an excellent medium for informal transmission of reports, recent results and bulletins to users with related interests and problems. The powerful screen-oriented editors available greatly enhance our capabilities for writing both text and programs.

Finally, the development of simulation programs generally requires a highly interactive computing environment - the sort of environment we feel is provided by SUMEX-AIM.

6.1.3 MOLGEN PROJECT

MOLGEN - An Experiment Planning System for Molecular Genetics

Prof. J. Lederberg (Genetics, Stanford)
Prof. N. Martin (Computer Science, U. of New Mexico)
Prof. E. Feigenbaum (Computer Science, Stanford)

I. Summary of Research Program

A. Technical Goals

The goal of the MOLGEN project is to develop an experiment planning system for the domain of molecular genetics. In order to accomplish this, we hope to create and apply innovative methods of knowledge management and hierarchical planning.

Experiments in molecular genetics are concerned with the study and manipulation of DNA molecules. The MOLGEN knowledge base will include both declarative and procedural information about such structures and the laboratory tools and techniques which experimental geneticists use. Also represented will be much of the strategic information required to join individual experimental steps into a meaningful whole. We are using the uniform method of schemata for representation of all types of knowledge within MOLGEN. We believe this will facilitate knowledge acquisition and explanation and provide a consistent means of storing hierarchical and other relations among objects and rules in the system. We hope to make the underlying knowledge base flexible enough to allow for experimentation with a wide variety of specific planning strategies.

B. Medical relevance and collaboration

Molecular genetics has at least two major connections to medical research. Learning about the basic mechanisms which control the operation and transmission of genetic information is necessary to understand and treat the wide range of diseases (and health conditions like aging) which are genetically controlled. Also, recent developments in molecular genetics offer the promise of using genetic mechanisms to produce essentially limitless amounts of drugs and other biomedical substances. The MOLGEN project will develop a system designed to aid the molecular geneticist in planning experiments of these types.

The MOLGEN project is a joint effort of the Computer Science Departments of Stanford and the University of New Mexico and the Genetics Department of Stanford. Major participants are Professor Nancy Martin of the University of New Mexico, Professor Edward Feigenbaum, Peter Friedland, Jonathan King, and Mark Stefik of Stanford Computer Science, and Professor Joshua Lederberg and Jerry Feitelson of Stanford Genetics.

C. Accomplishments

MOLGEN is in the first year of formal funding as an independent entity. We have devoted this year to learning and analyzing the basic knowledge of experimental molecular genetics and to building part of the central structure of the knowledge base management system. A wide variety of experiments have been studied with the aim of extracting knowledge about the genetic objects and operators used as well as the higher-level knowledge used to form the overall experimental plan. The object level knowledge is currently being organized into the schemata formalism for an initial attempt at a molecular genetics knowledge base.

A representation method for DNA structures and an interactive structure editing and entry system (EDNA) has been built and tested successfully with geneticist users. Work is proceeding on the schemata storage and access routines and on routines for acquiring and editing the rules which describe the procedural knowledge of the domain. We plan to have the basic MOLGEN system operational for the purpose of testing object and operator knowledge (the practical goal of experiment checking) by the end of July 1977.

D. Publications

- 1) N. Martin, P. Friedland, J. King, M. Stefik, "Knowledge Base Management for Experiment Planning in Molecular Genetics," submitted to Fifth International Joint Conference on Artificial Intelligence
- 2) M. Stefik and N. Martin, "A Review of Knowledge Based Systems as a Basis for a Genetics Experiment Designing System," Feb. 1977 Stanford CS Report STAN-CS-77-596, HPP77-5
- 3) N. Martin, P. Friedland, M. Stefik, "MOLGEN Knowledge Base I: Object System" To appear as HPP Working Paper
- 4) N. Martin, P. Friedland, M. Stefik, "MOLGEN Knowledge Base II: Rule System" To appear as HPP Working Paper

II. Interactions with SUMEX-AIM

All system development has taken place on the SUMEX-AIM facility. We have used the system not only for programming, but also as a major aid in writing and transmitting among ourselves the wide variety of formal and informal reports which are necessary in the MOLGEN design phase. We believe the availability of good interactive text editing facilities like TV-Edit increases our productivity significantly.

Active collaboration with remote users at the University of New Mexico will begin in September 1977 (Prof. Nancy Martin has been visiting at Stanford this year). We expect this collaboration to occur over the ARPA network. We hope also to maintain a collaboration with Dusko Ehrlich, formerly a Stanford geneticist and now doing research at The Institut de Biologie Moleculaire Faculte de Science in Paris over a TYMNET link to Sumex.

We have benefited enormously from the collected expertise in both knowledge-based systems and general programming and design problems available from other SUMEX-AIM projects. We have especially strong ties to the knowledge management expertise of the MYCIN project, but we also share common objectives with parts of the DENDRAL, SECS, and protein crystallography projects. We have also benefited from the intense interaction with many other projects at the AIM conferences.

Finally, we have provided small amounts of SUMEX resources to geneticist users as part of a quid pro quo relationship for helping us understand that subset of genetic knowledge necessary for our initial knowledge base. The most outstanding example of this sort of collaboration occurred with Prof. Larry Kedes' group at the VA hospital in Palo Alto who are using SUMEX to determine the feasibility of automated assistance in analyzing complex DNA base sequences.

6.1.4 MYCIN PROJECT

MYCIN - Computer-based Consultation in Clinical Therapeutics

S. N. Cohen, M.D. (Pharmacology) and
B. G. Buchanan, Ph.D. (Computer Science)
Stanford University

I) Summary of research

Technical goals

The Mycin project is aimed at the development of a computer program capable of functioning as an expert consultant on a range of medical decision making problems. In particular, we have been working on the construction of a system that provides consultative advice on the diagnosis and therapy selection for a number of infectious diseases. Current areas of competence of the system include bacteremia and meningitis, and work is currently underway to extend this to urinary tract infections, pulmonary infections, and prophylactic use of antibiotics.

Our work has been guided by three fundamental objectives:

- (1) A major goal of the MYCIN system has been to provide a computer-based therapeutic tool designed to be clinically useful, one that would be used eventually in the clinical setting. This goal requires development of a system that has a medically sound knowledge base, and that displays a high level of clinical competence in its field. The program must first convince clinicians of the quality of the information it is providing before they will be willing to use it.
- (2) Since many clinicians are not likely to accept the advice provided by a computer-based system unless they can understand why the recommended therapy has been selected, the system has to do more than just give advice dogmatically. It should have the ability to explain the reasoning behind its decisions, and should be able to do so in terms that suggest to the physician that the program approaches the problem in much the same way that he does. This permits the user to validate the program's reasoning, and modify (or reject) the advice if he believes that some step in the decision process is not justified. It also gives the program an inherent instructional capability that allows the physician to learn from each consultation session.
- (3) A third major goal is to provide the program with capabilities that enable augmentation or modification of the knowledge base by clinical experts in infectious disease therapy, in order to improve the validity of future consultations. The system therefore requires some capability for acquiring knowledge by interacting with experts in the field, and for incorporating this knowledge into its knowledge base.

Three separate parts of the MYCIN system accomplish these goals. The consultation system uses the knowledge base, along with patient-related data entered by the physician to generate therapeutic advice. The explanation system has the ability to explain the reasoning used during the consultation, and to document the motivation for questions asked or the rationale for conclusions reached. Finally, the knowledge acquisition system enables experts in antimicrobial therapy to update MYCIN's knowledge base, without requiring that they know how to program a computer.

We have also sought to use Mycin as a framework for understanding the process of medical decision making and the nature of clinical judgment. Physicians are constantly faced with the necessity of making decisions based on information that is both incomplete (missing historical data or test results not yet available) and inexact (results are rarely definitive). In addition, those decisions are often based on rules that are only approximate (e.g., "a gram-negative aerobic rod in the blood is probably a bacteriodes"). But decisions are made despite these problems, and the results often proven later to be valid. We have attempted to understand how this is done by developing in our system a parallel set of capabilities. We have relied on the "production rule" encoding of information, in which individual decision rules are specified in an "if/then" format. For example, the rule indicated just above is encoded in the system as:

```
If      1) the gram stain of the organism is gram negative, and
        2) the morphology of the organism is rod, and
        3) the aerobicity of the organism is anaerobic,
Then there is suggestive evidence (.6) that the identity of the organism
    is Bacteroides.
```

This encoding of knowledge offers a number of advantages over some of the more traditional approaches to diagnosis like decision trees, Bayesian analysis, and utility theory. Unlike decision trees, it can deal with both inexact and incomplete information. Unlike the Bayesian and utility theory approaches, it does not need extensive amounts of conditional probability data. A collection of independent rules is also far easier to augment than a complex decision tree; the rules thus provide a much more flexible body of knowledge to which new information is more easily added. The rules also make possible an explanatory capability: the system can justify any of its actions or decisions by displaying the relevant rules it invoked in reaching that decision. This provides an explanation that is far more comprehensible than any we might be able to provide by recapping the actions of a program based solely on statistical considerations.

A more specific goal of our research involves understanding the process of infectious disease diagnosis and therapy selection. This process is not as yet well understood, and we believe that by dissecting it down to individual decision rules, we can gain insight into how it works. In addition, the resulting set of rules may prove to be a useful compendium of knowledge about the task.

Since we believe this set of rules will also be quite large, we are studying the problems of accumulating, managing, and using large stores of such task-specific knowledge. We are working on a range of techniques to provide capabilities like insuring the consistency of the set of rules and making it easy to modify existing rules or add new ones.

Finally, since computer consultants are designed for use by people who might not otherwise make use of computers, we have devoted a great deal of attention to the issue of human engineering, and the "habitability" of the system. This ranges from such minor items as the automatic correction of misspelled answers, to the range of sophisticated explanation capabilities available.

Medical relevance and collaboration

A number of recent studies indicate a major need to improve the quality of antimicrobial therapy. Almost one-half of the total cost of drugs spent in treating hospitalized patients is spent on antibiotics [1,2], and if results of a number of recent studies are to be believed, a significant part of this therapy is associated with serious misuse [2,3,4,5]. Some of the inappropriate therapy involves incorrect selection of a therapeutic regimen [4], while another serious problem is the incorrect decision to administer any antibiotic [2,4,5]. One recent study concluded that one out of every four people in the United States was given penicillin during a recent year, and nearly 90% of these prescriptions were unnecessary [6]. Other studies have shown that physicians will often reach therapeutic decisions that differ significantly from the decisions that would have been suggested by experts in infectious disease therapy practicing at the same institution.

Nonexperts sometimes choose a drug regimen designed to cover for all possibilities, prescribing either several drugs or one of the so-called "broad spectrum" antibiotics, even though appropriate use of clinical data might have led to more rational and less toxic therapy. Within a hospital environment in which professional resources are often overburdened, and in environments where expert sources are not readily available, a computer-based consultant will be highly useful. Such a system will also have broad fringe benefits in its educational impact on staff physicians and in providing a framework for quality control and peer-review evaluations.

Antimicrobial therapy appears to be an especially suitable area for the initial development of a computer-based system to assist physicians with decisions in clinical therapeutics. The components of the decision making process in antimicrobial therapy are more readily definable than in many other areas of medicine, and the consequences of the physician's decision can usually be assessed in terms of the direct therapeutic action. Nevertheless, the general approach used here is applicable to other areas of clinical decision making. The basis of rational antimicrobial therapy decisions is identification of the microorganisms causing the infectious disease. Accurate identification is important because of the specificity of antibiotic action: drugs that are highly effective against certain organisms are often useless against others. The patient's clinical status and history (including information such as prior infections and treatments) provide data that may be valuable to the physician in identifying the disease-causing organisms. However, bacteriological cultures that use specimens taken from the site of the patient's infection usually provide the most definitive identifying information.

Initial culture reports from a microbiological laboratory may become available within 12 hours from the time a clinical specimen is obtained from the

patient. While the information in these early reports often serves to classify the organism in general terms, it does not often permit precise identification. It may be clinically unwise to postpone therapy until such identification can be made with certainty, a process that usually requires 24 to 48 hours, or longer. Thus it is commonly necessary for the physician to estimate the range of possible infecting organisms, and to start appropriate therapy even before the laboratory is able to identify the offending organism and its antibiotic sensitivities. In this setting MYCIN plays two roles: (a) providing consultative advice that will assist the physician in making the best therapeutic decision that can be made on the basis of available information, and, (b) by its questioning of the physician, pinpointing the items of clinical data that are necessary to increase the validity of the clinical decision.

Our project is an interdisciplinary effort involving the joint effort of computer scientists from the Stanford Computer Science Department, and clinicians from both the Department of Clinical Pharmacology at Stanford and the Department of Infectious Disease at the University of Arizona. The task of the clinicians has been to specify the decision rules necessary for diagnosis and therapy selection, while the computer scientists have been devising ways to represent and use this information in the computer. The system is then tested by the clinicians using real cases obtained from journals and medical records.

A complete listing of the staff is given below.

Stanley N. Cohen, MD, Clinical Pharmacology
Bruce G. Buchanan, PhD, Computer Science
Stanton Axline, MD, Infectious Disease (now at University of Arizona)
Randall Davis, PhD, Computer Science
Frank Rhame, MD, (to 9/75), Infectious Disease
Edward Shortliffe, MD PhD (to 6/76, returning 6/77), Infectious Disease
Victor Yu, MD, Infectious Disease
Rudolpho Chavez-Pardo, MD, (to 9/75), Clinical Pharmacology
A. Carlisle Scott, MS, Computer Science
Sharon Wraith, BS, Clinical Pharmacology
Jan Aikins, BS, Computer Science
Robert Blum, MD, presently in Computer Science
William Clancey, AB, Computer Science
Larry Fagan, AB, Computer Science
William van Melle, AB, Computer Science

Progress Report

Period covered: June 1, 1974 through September 30, 1976

Summary

Over the past three years we have designed, built and partially evaluated a computer program capable of diagnosis and therapy selection for certain varieties of infectious diseases. The program is intended to function as a consultant, and "interviews" a doctor about his patient, requesting information on clinical findings and results of laboratory tests. It relies on a store of judgmental knowledge (obtained from experts in infectious disease) to determine the

conclusions which can be drawn from the answers it receives. This judgmental knowledge is in the form of some 400 decision rules dealing with the wide range of topics that must be considered in determining the likely identity of causative organisms and selecting appropriate antimicrobials.

MYCIN is composed of the three systems described earlier (the consultation, explanation, and knowledge acquisition systems), all of which reference the knowledge base of decision rules. The program is currently capable of dealing with bacteremia and meningitis infections. It can diagnose the likely presence of more than 35 different organisms and can recommend therapy for 100 organisms, selecting drugs from a "pharmacopoeia" of 30 antimicrobials. The system can tailor its therapy recommendations to a specific organism and infection, can adjust dosage levels and durations in response to impaired renal status, and can combine drugs to create combination therapies, giving it a wide range of clinical applicability.

Detailed Report

Our work in the past several years has been organized around five main areas of investigation. We have

- a) increased the system's competence in existing areas of clinical expertise while expanding its scope
- b) developed a number of user-oriented features to increase the program's attractiveness to clinicians
- c) developed a range of knowledge acquisition capabilities to speed the process of expanding the system's clinical competence
- d) solved a number of technical problems to insure that the program does not outgrow the computer resources available to it
- e) evaluated the system's level of expertise.

Clinical Capabilities

Since the primary qualification for any clinical consultant is competence in the domain, we have devoted significant effort to expanding MYCIN's knowledge base and widening its scope of competence.

For instance, the system was directed initially at patients with positive blood cultures, the basic methodology was generalized to support a much broader approach to the problem. MYCIN has now gained the ability to deal with infections from which the causative pathogen hasn't been isolated (e.g., pneumonia), or which haven't even been cultured (e.g., brain abscess). With this broadening of scope, it has also become necessary to be able to evaluate the meaningfulness of isolates for cultures taken from sites other than blood. For urine and sputum isolates, for example, the system gained the ability to base its evaluation of sterility of an isolate on both the method of collection and the user's estimation of conscientiousness of collection.

An extensive review of the program's approach to drug selection has led to a major revision in the basis for therapy selection during the course of program development. The program was given the ability to consider both the infectious disease diagnosis and the significance of the organism as further determinants of therapy, in addition to organism identity. These three together have become the primary factors in drug selection, with drug toxicity and ecological factors as secondary considerations. The result is a more appropriate, more sharply focussed drug selection that also includes dose, route, and duration.

While the initial development of the knowledge base focussed on rules concerned with the diagnosis and therapy for blood infections (bacteremia), the complexity of infectious disease therapy and the frequent occurrence of multiple infections in a single patient requires a broader knowledge if the system is to be clinically useful.

In response we have extended MYCIN's knowledge base, while at the same time improving the degree of sophistication with which the system deals with bacteremia. The second major area has been the diagnosis and treatment of meningitis, and more than 100 rules were added to provide the ability to deal with it. In the process the program was also extended beyond bacteria, as it gained the ability to consider and treat both fungi and viruses.

This area has proved to be an especially useful domain because it has presented several new challenges. In particular, meningitis requires the ability to deal with a disease that is often diagnosed on clinical grounds alone, before any specific microbiological evidence is available (by comparison, the diagnosis of bacteremia on clinical grounds alone is far less certain, and usually requires establishment of the fact that bacterial growth has occurred in blood cultures.) For this reason, extension of the project into the meningitis area has made it necessary for MYCIN to consider a larger range of clinical factors, and has resulted in a system which has a broader picture of the whole patient.

Other contributions to the system's competence have come from expansion of the knowledge base to include information about normal bacteriological flora for a wide range of culture sites. This enables the program to distinguish between normal and pathological flora, and it can as a result decide more precisely on whether to treat.

User Oriented Features

Clinicians traditionally shun computer programs, and we believe this is in large measure due to insufficient attention paid to user oriented features. As a result, we have devoted significant effort to insuring that MYCIN is responsive to its users in a number of unique ways. The development of the explanation and question answering capabilities have been a essential for this work, and both have grown extensively in power.

The system's ability to explain the motivations for its questions, for instance, underwent a major design revision. It is now based on a more powerful approach that relies on the program's knowledge of its own control structure and ability to examine its own rules. The user can now fully explore the system's current line of reasoning, rather than just a single level, as initially implemented.

The language understanding capabilities of the question answering system have also been extensively revised. They now allow a broader range of questions to be asked and offer more precise answers. The use of this feature was also simplified so that the user no longer needs to classify his questions.

A comprehensive review of the kinds of questions asked by users of the system has led to a number of important features. MYCIN can now answer a much wider range of questions, and can, in particular, explain why it did not take a specific action, as well as why positive conclusions were reached. It is our feeling that capabilities such as these are of great importance in enabling the project's staff and clinical experts to understand the program's rationale for its actions in instances where its recommendations do not appear to be the most appropriate and most correct. Thus, the line of reasoning of the program can be evaluated, and requirements for new or modified rules can be uncovered. These kinds of capabilities are also important in optimizing user acceptance of the system.

A substantial addition to the question-answering facility enables the system to explain the process of therapy selection. In comparison to the diagnostic process, therapy selection is complicated somewhat by the need to consider a range of different factors simultaneously, such as the total number of drugs recommended, the degree of sickness of the patient, possible interactions between drugs, toxicity and other side effects, etc. Despite this complexity, explanations of therapy selection are phrased at a conceptual level that makes them comprehensible to the physician. As before, this makes it possible for the physician to verify the validity of the system's decisions, and makes it clear to him that the system reaches its results in much the same way that he does.

The explanation consists of a step-by-step review of the reasoning which led to recommending a particular drug for a specific organism. It considers such issues as why a drug was first considered for an organism, why a drug may have been chosen as the best therapy for that organism, how the total number of drugs was reduced by considering common drug classes among the candidates, and consideration of possible contraindications based on the patient's allergies, age, and other factors. By characterizing each drug according to this scheme, the program can explain why a drug was or wasn't prescribed, as well as why one drug is to be preferred over another. This offers an important explanatory capability that will make the system more attractive and acceptable to clinicians.

Several capabilities have been added to make the program easy to use. The system is now more tolerant of erroneous or inappropriate responses, and is able to provide a reworded question, along with a list of acceptable answers. In addition, it has the ability to recognize responses which are not sufficiently precise, and can rephrase its questions accordingly.

We have recently added to the system the ability to modify drug dosage in cases of renal failure. Where, previously, the system only issued a warning to modify doses, it is now able to use either creatinine clearance or serum creatinine levels to compute the level of renal function. The program then uses drug-specific information (e.g., half-life, percent loss of the drug via renal excretion, etc.) to adjust the regimen. It can either (a) adjust dose levels downward and leave dosing interval unchanged, or (b) increase dosing interval and

leave levels unchanged, or (c) allow the physician to select a dose interval, for which it chooses an appropriate dose level.

Since the problem of determining renal status and the proper adjustment of drug dose is important in the use of aminoglycoside antibiotics, cephalosporins, and other antimicrobial agents, the customization of drug dosage recommendations will be an important addition to the power of the system.

We have found, in addition, that there is a substantial amount of information that is routinely collected in every consultation, like the date and site of each of the cultures, gramstain and morphology results for each of the organisms that grew out, etc. Currently, the program exhaustively analyzes each culture and all of its organisms in turn. Some users of the program appear to be impatient with this method, and would much prefer to enter all the relevant data on all the cultures and organisms at once. This is faster and easier, since the information can be gathered in a single review of the chart, instead of having to review it several times as each culture is processed. In response to this, we have reorganized the consultation slightly, so that it is possible to enter all of this data at once, at the beginning. This offers two other advantages in addition to improving the program's acceptability to its users. First, it provides a basis for our future efforts to write rules which deal with interactions between infections (see below, "Specific Aims"), and second, it suggests a mechanism for eventually merging our work with the product of existing efforts to organize and automate the recording and handling of medical record data. This latter development may in time make it possible for MYCIN to obtain a large part of the information it requires directly from such automated records, sharply reducing the number of questions it has to ask, and speeding up the consultation considerably.

Finally, several new capabilities make the system convenient to use, in anticipation of its evaluation in the clinical setting. Among these are the option of the user to type a comment about system performance at any time during the consultation. His comment is recorded in a special file which is reviewed periodically by our medical staff, and provides an on-going opportunity for users to offer feedback aimed at improving the usefulness of the system. The user can also indicate his belief that the system has "broken down" in some way and he is invited to describe the problem. His description is saved along with information about the current state of the program, so that our systems programmers can deal with the problem later.

Knowledge Acquisition

A preliminary knowledge acquisition program was completed in the middle of 1974, and demonstrated the feasibility of having a physician teach the system new rules using a rather stylized subset of English. Building on the experience gained here, work began on a revised program designed to allow the user to examine and modify the program's knowledge and behavior as a single, unified action. This program was designed to make the explanation and knowledge acquisition capabilities available together, to make use of the fact that the nature of the explanations requested can give a clear hint about the content of a new rule. The program was also designed to advise the user about the effect of his rule on the original deficiency, indicating, for instance, whether or not it corrects the problem he noticed.

Work on a preliminary version of this new program was completed in 1976, making available a broad range of useful features enabling our clinical experts to add rules to the system without requiring that they have a knowledge of programming. If the expert finds that MYCIN's handling of a particular problem is at variance with his own expert knowledge, he can use the explanation capabilities to discuss the line of reasoning in use at that time, can add or modify rules in the knowledge base, and can determine the effects of the changes on MYCIN's subsequent performance. (Quality control is maintained on the overall system by regular meetings of our clinical and pharmacological experts who determine the "official" MYCIN knowledge base.)

Technical Issues

As MYCIN's clinical capabilities have expanded, efficiency has improved as a result of a number of modifications to the system's technical capabilities. Early in our work, for instance, a comprehensive review and modification of the control structure was undertaken to improve efficiency and generality. The resulting program was both more direct, and faster.

More recently, modifications have been made so that the the large English dictionary can be kept on the disk and accessed only as needed, rather than keeping it in core, which slows down the system's response speed. The self documenting features of the program have also been improved to make them faster, and the system's interaction with the terminal has been made more uniform, to prepare for the time when different users of the system may have various different kinds of terminals.

Evaluation Activities

Since clinicians are likely to require documentation of MYCIN's competence and utility before seeking its advice, considerable time has been spent on evaluating the system and on implementing a range of program features to support these efforts.

In the past two years we have obtained many useful suggestions from clinicians when the system was presented to several different conferences. In February 1975 it was presented to the Western Society for Clinical Research, in September 1975 to the International Symposium on Clinical Pharmacy and Clinical Pharmacology, and more recently (June 1976), it was presented to the Drug Information Association.

A large scale formal study and evaluation of MYCIN's performance was begun in January 1976. The same set of clinical data was provided to both MYCIN and a set of experts in infectious disease therapy. [Five of the experts were nationally recognized authorities in the field, the other five were clinical fellows in the Infectious Disease Division at Stanford. A complete list of names, titles and affiliations is found in Appendix B.] The judgments of the program and the experts were compared, and the experts were asked to evaluate MYCIN's performance.

To do this, we first designed a form to allow us to separate the variables requiring analysis. The parameters evaluated include

- A. the "quality" of the interaction - were any questions irrelevant or missing
- B. the program's ability to determine organism identity
- C. the program's ability to determine organism significance
- D. the program's ability to select proper therapy
- E. overall performance evaluation
- F. potential impact as a clinical tool or teaching facility

The evaluation form was designed to be informative yet simple to complete. It was tested in a pre-evaluation trial run, then used for the formal study.

Consecutive patients with positive blood samples were evaluated for inclusion in the study by project personnel, until we obtained at least 10 patients for which MYCIN recommended therapy, and 15 patients overall (patients were rejected if they were outpatients when the sample was drawn, if they had a previous blood culture in the preceding seven days, or if they had a diagnosis of meningitis or infectious endocarditis.) For each of the patients accepted, a one to two page clinical summary was prepared and combined with a summary of the laboratory test data as of the time when the first blood culture was obtained. This information was then used to obtain a therapeutic evaluation from MYCIN.

Each of the participating experts received a set of fifteen evaluation forms (one for each patient). Each form contained: (a) the clinical summary and lab data; (b) space for the expert to record his conclusions about the nature of the infection, likely causative organisms, and appropriate therapy; and (c) a transcript of the MYCIN consultation along with space for the expert to record his opinion of various aspects of MYCIN's performance. By presenting the information in this order, we obtained a therapeutic regimen from the expert based on the same information supplied to MYCIN. This allowed us to compare the expert's answers to MYCIN's, and also gave us the expert's opinion of the system's performance.

In the past few months a sufficient number of the forms have been returned that we were able to do a preliminary analysis. The figures below are based on the nine (out of ten) which have been returned.

Since it is difficult to select a single number which summarizes performance, we have in general measured each of the parameters listed above in three ways: (i) the percent of instances in which the program was judged exactly correct, (ii) the percent of instances in which the program's performance was judged exactly correct or an acceptable alternative, and (iii) the percent of cases in which a majority of the experts judged its performance exactly correct or an acceptable alternative. By using all three measures, we obtain a range of figures which give a good picture of the program's performance.

All of these attempts to evaluate performance are complicated by the fact that (as expected) the experts' own choices about each patient were not unanimous. Thus, we cannot ask whether MYCIN's answers were "correct" in any absolute sense, since there was no agreement on what constitutes "correct". Instead, we ask how often each individual expert rated the program's responses as

correct. But given the variation among experts themselves, the program can never be expected to reach 100%, and depending on the extent of the intra-group variation, the absolute limit may in fact be much lower. Thus the ideal question to ask is "Do experts rate MYCIN's performance correct at least as often as they rate each other's performance correct?" This would give a good indication of how close the system's performance was to that of the group of experts as a whole.

We have been able to do this in a few isolated cases, but in general it requires more information than we were able to collect. This is discussed in more detail below, but in general terms the problem is that we were able to ask each expert for his choices for each patient, and ask him to rate MYCIN's choices. But, without a second round of questionnaires, which would ask each expert to rate the acceptability of the other 9 experts' responses, we lack direct information about intra-expert variability. The figures below should be reviewed with this caveat in mind.

A. "Quality" of the interaction

To measure the first item, the experts were instructed to mark any questions in the consultation which they felt were irrelevant, and to note any questions which they felt were omitted by the system. Overall MYCIN did quite well, as there were no consultations in which a majority of the experts felt that any particular question was irrelevant or omitted. On the average, there were 0.53 questions judged irrelevant and 0.55 indicated as omitted.

Table I summarizes the next four measurements.

	MYCIN 1st choice identical to an expert's 1st choice	MYCIN 1st choice identical to or an acceptable alternative to an expert's 1st choice	MYCIN 1st choice identical to or an acceptable alternative judged by a majority of experts
ORGANISM IDENTITY	56.3% N= 414	75.6% N= 414	81.8% N= 11
ORGANISM SIGNIFICANCE	91.7% N= 36	NA	100% N= 4
THERAPY SELECTION	12% N= 99	75% N= 99	91% N= 11
OVERALL PERFORMANCE	17.0% N= 135	59.3% N= 135	60.0% N= 15

Table I
Summary of nine experts' responses to MYCIN's performance on 15 cases

B. Organism Identity

For organism identity, the experts were asked to rate each of MYCIN's selections as exactly correct (they agreed that the organism was likely to be present), an acceptable alternative (they had not chosen that organism, but agreed it might be present), or an unacceptable choice (they disagreed with its selection). Since 11 of the cases were not contaminants, and there was a total of 46 organisms chosen by the system, with 9 experts rating each of those choices we have an N of 414 for the first two columns and 11 for the third.

In 56% of the instances the system's choices were identical to the experts', 75% of them were either identical or acceptable alternatives, and in 82% of the cases, its results were acceptable to a majority of the experts.

In addition, the experts were asked to indicate which organisms they felt MYCIN had overlooked in its diagnosis. For the 11 non-contaminant cases, the experts indicated an average of only 0.35 organism identities that were overlooked by the system. In no case did a majority of experts feel that any particular organism had been overlooked, suggesting that even the 0.35 figure is a result of intra-expert variation.

C. Organism Significance

The first question on the evaluation form gave the expert a chance to indicate that he felt the patient did not need to be treated. The first column of the second row indicates the number of times the expert indicated no treatment was necessary for a case in which MYCIN also judged the organism to be a contaminant. (There is no number in the second column since we did not ask about a "close call" on whether or not to treat. In addition, the measurement is based only on the contaminant cases, since in many of the cases where both MYCIN and the expert determined that treatment was necessary, they based that decision on different organisms. We felt that it would be misrepresentative to call these situations "agreements".)

As the figures show, in only three out of 36 instances was there any disagreement with the system's decision on whether or not to treat.

D. Therapy Selection

The expert was asked to select therapy for the organisms which he felt were likely to be present before looking at MYCIN's therapy recommendation. He was then asked to judge MYCIN's choice of therapy for that patient. Since MYCIN was selecting therapy for the organisms which it felt were present (which may have differed from those chosen by the expert), this provides a fundamental comparison of performance - it compares therapy selection performance of the two when they are faced with the same clinical situation.

This comparison is a difficult one to make, since it is complicated by the difficulty noted above, of variability in the experts' performance and the need to judge MYCIN with respect to that variability. Looking only at exact agreements (i.e., two identical therapies) produces the figure in the first column, which indicates that 12% of the time MYCIN's recommendation was identical to that of an expert. Comparing each expert's therapy choice with the other 8 indicates that 35% of the time (N= 396) any pair of experts chose identical regimens. The experts were also asked to judge whether MYCIN's therapy was an acceptable alternative (if it was not identical to their own), producing the figure in the second column. This indicates that it was either identical, or they felt it was an acceptable alternative 75% of the time. (Unfortunately, we have no reliable way of judging the intra-expert variability here, without a second round of questionnaires which asked each expert to rate the acceptability of the other experts' choices.) [As an alternative, we have attempted to develop a measure of how "far apart" two non-identical regimens are. But the problem is difficult: for example, for gram negative rods with salmonella most likely, is gentamycin and chloramphenicol "very different" from gentamycin and ampicillin? We have been working on a "drug metric" to solve this problem, attempting to base the "difference" between two drugs on factors like organism susceptibility, toxicity, and drug efficacy, but this work is still in progress.]

The figure in the third column gives a crude overall measure of therapy selection performance, and indicates that in 91% (10 out of 11 cases), a majority of the experts rated MYCIN's regimen as either identical to their own or an acceptable alternative.

[The evaluation form also asked each expert to choose a regimen for the organisms which MYCIN had selected. The intent here was to compare the system's performance against the expert when both were faced with the same set of organisms (rather than compared with the same clinical situation, as above). Unfortunately, inconsistent answers on the part of the experts indicated that they were not answering the question according to the instructions. It appeared that they were not able to suspend their own judgments about organism identity sufficiently to select a regimen based on MYCIN's organisms alone. For this reason, we believe the data to be unreliable, and have not included it here.]

E. Overall Performance

At the end of each evaluation form, the expert was asked to rate the system's overall performance as either excellent, good, fair, or poor. The first two columns of the last row indicate that 17% of these evaluations were "excellent", and almost 60% were either "excellent" or "good" (only 13% were "poor"). In 60% of the cases (9 out of 15), a majority of the experts felt that MYCIN's overall performance was either "excellent" or "good".

F. Present Utility and Future Potential

Finally, after completing the entire set of 15 patients, each expert was asked to rate MYCIN's present utility and future potential as a clinical tool and as an educational tool, rating it as having "considerable", "some", or "no" potential. The table below summarizes their response.

Evaluation of Present Utility

	"considerable"	"some"	"none"
clinical tool	11%	67%	22%
educational tool	11%	89%	0%

Evaluation of Future Potential

	"considerable"	"some"	"none"
clinical tool	11%	89%	0%
educational tool	67%	33%	0%

Table II

Opinions of 9 experts on MYCIN's present utility and future potential

To aid these evaluation efforts, we have also implemented a number of useful features in the system. For instance, MYCIN now keeps continuing

statistics of the use of rules in its knowledge base. This will help us to monitor its long term performance, to study the interrelationship between rules, and perhaps detect automatically any inconsistencies or gaps in the knowledge base.

We have also designed and implemented a mechanism for "on-line" evaluation. At the end of each consultation, the system asks a few questions about the quality of its performance from the clinicians who are using it. This interchange will be brief to avoid being a burden to the user, but it is expected to represent an important addition to the other evaluation efforts.

It will, for instance, make possible a new form of evaluation of the system. Rather than using a series of "prepackaged" cases as was done in our initial evaluation, the next stage will be carried out using information entered at a terminal by the evaluator. The participating panel of experts will be selecting patients in areas covered by the MYCIN knowledge base, and will engage in a dialogue with the system about those patients. Following completion of the session, the on-line evaluation feature will ask questions about system performance, and the responses will be tabulated and evaluated on-line by appropriate biostatistical programs. Specific recommendations which may point out problem areas in the consultation will be reviewed by our staff. By this process we expect to be able to maintain a continuing evaluation of MYCIN's capabilities in various areas, and pinpoint specific areas where performance is suboptimal.

MYCIN Project Publications

THESES

Davis R,

Applications of meta level knowledge to the construction, maintenance, and use of large knowledge bases, Thesis: PhD in Computer Science, AI Memo 283, 304 pp, Stanford University, July 1976.

Shortliffe E H,

MYCIN: A rule-based computer program for advising physicians regarding Antimicrobial therapy selection, Thesis: Ph.D. in Medical Information Sciences, Stanford University, Stanford CA, 409 pages, October 1974. Also, Computer-Based Medical Consultations: MYCIN, American Elsevier, New York, 1976.

PAPERS

Buchanan B G, Davis R, Yu V, Cohen S N,

Rule-based medical decision making by computer, Proc. MEDINFO 1977, to appear.

Clancey W.

Chronicler: an explanation system based on set-predicate representation of computational processes, submitted to 5th IJCAI.

Aikins J S.

Use of models in a rule-based consultation system, short paper submitted to 5th IJCAI.

Davis R.

Interactive transfer of expertise: acquisition of new inference rules, submitted to 5th IJCAI.

Davis R.

Knowledge acquisition in rule-based systems: knowledge about representations as a basis for system construction and maintenance, to appear in Pattern Directed Inference Systems, Waterman and Hayes-Roth (eds.), Academic Press, in press. Also to be presented at Pattern Directed Inference Systems Workshop, Honolulu, May 1977.

Davis R, Buchanan B G.

Meta-level knowledge: overview and applications, submitted to 5th IJCAI, Cambridge, MA, August 1977.

Davis R.

A decision support system for medical diagnosis and therapy selection, Data Base (SIGBDP newsletter), 8 (Winter 1977) pp 58-72.

Wraith S, Aikins J, Buchanan B G, Clancy W, Davis R, Fagan L, Scott A C, van Melle W, Yu V, Axline S, Cohen S,

Computerized consultation system for selection of antimicrobial therapy, American Journal of Hospital Pharmacy, 33 (December 1976) pp 1304-1308

Scott A C, Clancey W, Davis R, Shortliffe E H,

Explanation capabilities of knowledge based production systems, American Journal of Computational Linguistics, Microfiche 62, 1977. Also, HPP Memo 77-1, Stanford Computer Science Department, February 1977.

Shortliffe E H, Davis R,

Some considerations for the implementation of knowledge-based expert systems, SIGART Newsletter, 55:9-12, December 1975.

Davis R, Buchanan B, Shortliffe E H,

Production rules as a representation for a knowledge-based consultation system, Artificial Intelligence, 8 (Spring 1977) pp 15-45. (Also, AI Memo 266, Stanford University, October 1975).

Davis R, King J J,

An overview of production systems, in Elcock and Michie (Eds.), Machine Intelligence 8: Machine Representations of Knowledge, John Wyllie, to appear, 1977. (Also AI Memo 271, Stanford University, October 1975).

Shortliffe E H,

Judgmental knowledge as a basis for computer-assisted clinical decision making, Proceedings of the 1975 International Conference on Cybernetics and Society, pp 256-7, September 1975.

- Snortliffe E H, Axline S, Buchanan B G, Davis R, Cohen S,
A computer-based approach to the promotion of rational clinical use of
antimicrobials, in Gouveia, Tognoni and Van der Kleijn (Eds.), Clinical
Pharmacy and Clinical Pharmacology, pp 259-274, Elseiver/North Holland
Biomedical Press, 1976.
- E H Shortliffe, R Davis, S G Axline, B G Buchanan, C C Green, S N Cohen,
Computer-based consultations in clinical therapeutics: explanation and rule
acquisition capabilities of the MYCIN system, Computers and Biomedical
Research, 8:303-320 (August 1975).
- E H Shortliffe and B G Buchanan,
A Model of Inexact Reasoning in Medicine, Mathematical Biosciences
23:351-379, 1975.
- Shortliffe E H, Rhame F S, Axline S G, Cohen S N, Buchanan B G, Davis R, Scott A
C, Chavez-Pardo R, and van Melle W J
MYCIN: A computer program providing antimicrobial therapy recommendations
(abstract only). Presented at the 28th Annual Meeting, Western Society For
Clinical Research, Carmel, CA, 6 Feb 1975. Clin. Res. 23:107a (1975).
Reproduced in Clinical Medicine, p. 34, August 1975.
- Shortliffe E H
MYCIN: A rule-based computer program for advising physicians regarding
antimicrobial therapy selection (abstract only); Proceedings of the ACM
National Congress (SIGBIO Session), p. 739, November 1974. Reproduced in
Computing Reviews 16:331 (1975).
- E H Shortliffe, S G Axline, B G Buchanan, S N Cohen,
Design considerations for a program to provide consultations in clinical
therapeutics, Presented at San Diego Biomedical Symposium 1974 (February
6-8, 1974).
- E H Shortliffe, S G Axline, B G Buchanan, T C Merigan, S N Cohen.
An artificial intelligence program to advise physicians regarding
antimicrobial therapy, Computers and Biomedical Research, 6:544-560
(1973).

Articles About MYCIN

"Which Antibiotic?" Emergency Medicine, January 1977, pp 152-162.

II) Interactions with Sumex-Aim resource

Collaborations and medical use of programs

Dr. Jon Heiser

We have been working with Dr. Jon Heiser of the Department of Psychiatry of the University of California at Irvine, in an effort to create a consultant for the use of psychoactive drugs. We began by creating a version of Mycin that had all of the infectious disease knowledge removed from it, and showed Dr. Heiser how to build up the required base of knowledge about the new field.

He has, with his students, developed a small, but functional system that demonstrates encouraging performance on the task. Work has now begun in earnest to extend the competence of this pilot system, to produce a consultant with a useful level of performance.

It is interesting to note that the explanation capabilities required no modification whatever, and worked in the new system exactly as designed for the original system, despite the change in domains.

INTERNIST Project

The Sumex computer has made possible a valuable interaction between researchers on the MYCIN project at Stanford University and those working on the INTERNIST project at the University of Pittsburgh. These researchers are studying the possible representations and uses for disease models in a medical diagnosis system. Both research groups have been able to run each others programs and to study the medical knowledge bases which are stored on the Sumex computer. Communication between project members has also been greatly facilitated through use of the Sumex system.

Stanford Infectious Disease Faculty

Dr. Victor Yu of our group has been actively soliciting the involvement of the Stanford ID faculty in the development and evaluation of Mycin. He recently presented the system to the faculty and fellows of the Department, and has been seeking ways to involve the system in the Department's educational activities. For instance, medical students under his supervision have used the system during their ID rotation, comparing its results and reasoning process with their own on problems encountered in patients on the wards.

The Pulmonary Function Facility

Members of the Mycin project have also been collaborating with Dr. John Osborn and his co-workers of the Presbyterian Hospital/Pacific Medical Center in San Francisco on the development of a program to interpret the results of standard pulmonary function tests. The program is designed to perform a range of tasks, including: identifying the need to repeat tests because of poor patient effort; identifying the need for additional information in order to make a more definitive diagnosis; reporting and explaining the reasons for primary and secondary diagnoses and severity of any disease state; identifying the relation between diagnosis and any referral diagnosis; and interpreting any change from previous tests, or limitations on the interpretation because of the test methodology and the patient effort.

Sharing with other projects

Groups at Rutgers University, the University of Pittsburgh, Rochester University, and the University of Virginia Medical School have all been involved in varying degrees with running Mycin and evaluating its performance. They have suggested to us improvements in its design, and stock of medical knowledge, and made useful contributions to its development.

In addition, we have made use of the programs developed at both Rutgers and Pittsburgh. The former has been instructive to us in its handling of dynamically changing situations, while the latter has helped us to develop our own ideas about the modelling and use of prototypical descriptions of disease states.

The Molgen group at Stanford has also profited from much of our experience in acquiring knowledge and building large knowledge bases. Several of their

techniques for accumulating knowledge about genetics are based on extensions to ideas first suggested in some of our work.

In all of these cases, the use of Sumex as a national resource has clearly been a critical factor in making possible this sort of interaction.

References

- [1] Reiman H H, D'ambola J, The use and cost of antimicrobials in hospitals, Arch Environ Health, 13:631-636 (1966).
- [2] Kunin C M, et.al., Use of antibiotics: a brief exposition of the problem and some tentative solutions, Anns Int Med, 79:555-560 (1973).
- [3] Sheckler W E, Bennett J V, Antibiotic usage in seven community hospitals, J Amer Med Assoc, 213:264-267 (1970).
- [4] Roberts A W, Visconti J A, The rational and irrational use of systemic antimicrobial drugs, Amer J Hosp Pharm, 29:828-834 (1972).
- [5] Simmons H E, Stolley P D, This is medical progress? Trends and consequences of antibiotic use in the United States, J Amer Med Assoc, 227:1023-1026 (1974).
- [6] Kagan B M, Fanin S L, Bardie F, Spotlight on antimicrobial agents, JAMA, 226:306-310 (1973).

6.1.5 PROTEIN STRUCTURE PROJECT

Protein Structure Modeling Project

Prof. J. Kraut and Dr. S. Freer (Chemistry, U. C. San Diego) and
Prof. E. Feigenbaum and Dr. R. Engelman (Computer Science, Stanford)

I. Summary of research program

A. Technical goals

The goals of the protein structure modeling project are to 1) identify critical tasks in protein structure elucidation which may benefit by the application of AI problem-solving techniques, and 2) design and implement programs to perform those tasks. We have identified two principal areas which have both practical and theoretical interest to both protein crystallographers and computer scientists working in AI. The first is the problem of interpreting a three-dimensional electron density map. The second is the problem of determining a plausible structure in the absence of phase information normally inferred from experimental isomorphous replacement data. Current emphasis is on the implementation of a program for interpreting electron density (e.d.) maps.

B. Medical relevance and collaboration

The biomedical relevance of protein crystallography has been well stated in a recent textbook on the subject (Blundell & Johnson, Protein Crystallography, Academic Press, 1976):

"Protein Crystallography is the application of the techniques of X-ray diffraction ... to crystals of one of the most important classes of biological molecules, the proteins. ... It is known that the diverse biological functions of these complex molecules are determined by and are dependent upon their three-dimensional structure and upon the ability of these structures to respond to other molecules by changes in shape. At the present time X-ray analysis of protein crystals forms the only method by which detailed structural information (in terms of the spatial coordinates of the atoms) may be obtained. The results of these analyses have provided firm structural evidence which, together with biochemical and chemical studies, immediately suggests proposals concerning the molecular basis of biological activity."

The project is a collaboration of computer scientists at Stanford University and crystallographers at the University of California at San Diego (under the direction of Prof. Joseph Kraut) and at Oak Ridge National Laboratories (Dr. Carroll Johnson).

C. Progress summary

During the past year we have been designing and implementing a system of programs for interpreting three-dimensional e.d. maps. Progress has been made by attacking the problem from two directions: working upward from the primary data (i.e. the array of e.d. values) to higher level symbolic abstractions, and working downward from the given amino acid sequence and other experimental information to generate candidate structures which can then be confirmed by the abstracted data.

In the "bottom-up" area of research we have developed and implemented programs for analyzing topological features of the skeletonized e.d. map in terms of protein structural elements (e.g., side chains, chain ends, bridges, etc.), for finding local maxima, and, recently for generating a critical point network, i.e. a three-dimensional spanning tree which connects all critical points (peaks, saddle points) found in the map.

In the "top-down" area we have designed and implemented, in INTERLISP, a structure inference program which generates structural hypotheses at several levels of detail. At present the program can infer, from the amino acid sequence and other chemical information, and the symbolic abstractions of the e.d. map, the location of heavy atoms, cofactors and chain ends. Those features provide footholds, i.e. islands of certainty, from which additional structure is inferred by extension. Work is currently in progress on identification of the main chain, disambiguation of multiply connected regions and classification of side chain regions.

The system under development is knowledge-based. Both the corpus of knowledge of the task domain and the problem-solving strategy knowledge are incorporated as production-like rules.

D. List of Publications

- 1) Robert S. Engelman and H. Penny Nii, "A Knowledge-Based System for the Interpretation of Protein X-Ray Crystallographic Data," Heuristic Programming Project Memo HPP-77-2, January, 1977. (Alternate identification: STAN-CS-77-589)
- 2) E.A. Feigenbaum, R.S. Engelman, C.K. Johnson, "A Correlation Between Crystallographic Computing and Artificial Intelligence," in Acta Crystallographica, A33:13, (1977). (Alternate identification: HPP-77-25)

II. Interaction with the SUMEX-AIM resource

A. Collaborations

The protein structure modeling project has been a collaborative effort since its inception, involving co-workers at Stanford and UCSD (and, more recently, at Oak Ridge). The SUMEX facility has provided a focus for the communication of knowledge, programs and data. Without the special facilities provided by SUMEX the research would be seriously impeded. Computer networking has been especially effective in facilitating the transfer of information. For example, the more traditional computational analyses of the UCSD crystallographic data are made at the CDC 7600 facility at Berkeley. As the processed data, specifically the e.d maps and their Fourier transforms, become available, they are transferred to SUMEX via the FTP facility of the ARPA net, with a minimum of fuss. (Unfortunately, other methods of data transfer are often necessary as well -- see below.) Programs developed at SUMEX, or transferred to SUMEX from other laboratories, are shared directly among the collaborators. Indeed, with some of the programs which have originated at UCSD and elsewhere, our off-campus collaborators frequently find it easier to use the SUMEX versions because of the interactive computing environment and ease of access. Advice, progress reports, new ideas, general information, etc. are communicated via the message and/or bulletin board facilities.

B. Interaction with other SUMEX-AIM projects

Our interactions with other SUMEX-AIM projects have been mostly in the form of personal contacts. We have strong ties to the DENDRAL, Meta-DENDRAL and MOLGEN projects and keep abreast of research in those areas on a regular basis through informal discussions. The SUMEX-AIM workshop in June, 1976 provided an excellent opportunity to survey all the projects in the community. Common research themes, e.g. knowledge-based systems, as well as alternate problem-solving methodologies were particularly valuable to share. (That workshop was very likely the most significant conference for applied AI to be held in 1976.)

6.2 NATIONAL AIM PROJECTS

The following group of projects is formally approved for access to the AIM aliquot of the SUMEX-AIM resource. Their access is based on review by the AIM Advisory Group and approval by the AIM Executive Committee.

6.2.1 ACQUISITION OF COGNITIVE PROCEDURES (ACT)

Acquisition of Cognitive Procedures (ACT)

Dr. John Anderson
Yale University

I. Summary of Research Program

A. Technical goals:

To develop a production system that will serve as an interpreter of the active portion of an associative network. To model a range of cognitive tasks including memory tasks, inferential reasoning, language processing, and problem solving. To develop an induction system capable of acquiring cognitive procedures with a special emphasis on language acquisition.

B. Medical relevance and collaboration:

1. The ACT model is a general model of cognition. It provides a useful model of the development of and performance of the sorts of decision making that occur in medicine.
2. The ACT model also represents basic work in AI. It is in part an attempt to develop a self-organizing intelligent system. As such it is relevant to the goal of development of intelligent artificial aids in medicine.

We have been evolving a collaborative relationship with Dr. James Greeno and Allan Lesgold at the University of Pittsburgh. They are applying ACT to modeling the acquisition of reading and problem solving skills. We plan to make ACT a guest system within SUMEX. ACT is currently at the state where it can be shipped to other INTERLISP facilities. We have received a number of inquiries about the ACT system. ACT is a system in a continual state of development but we periodically freeze versions of ACT which we maintain and make available to the national AI community.

C. Progress and accomplishments:

ACT provides a uniform set of theoretical mechanisms to model such aspects of human cognition as memory, inferential processes, language processing, and problem solving. ACT's knowledge base consists of two components, a propositional component and a procedural component. The propositional component is provided by an associative network encoding a set of facts known about the world. This provides the system's semantic memory. The procedural component

consists of a set of productions which operate on the associative network. ACT's production system is considerably different than many of the other currently available systems (e.g., Newell's PSG). These differences have been introduced in order to create a system that will operate on an associative network and in order to accurately model certain aspects of human cognition.

A small portion of the semantic network is active at any point in time. Productions can only inspect that portion of the network which is active at the particular time. This restriction to the active portion of the network provides a means to focus the ACT system in a large data base of facts. Activation can spread down network paths from active nodes to activate new nodes and links. To prevent activation from growing continuously there is a dampening process which periodically deactivates all but a select few nodes. The condition of a production specifies that certain features be true of the active portion of the network. The action of a production specifies that certain changes be made to the network. Each production can be conceived of as an independent "demon." Its purpose is to see if the network configuration specified in its condition is satisfied in the active portion. If it is, the production will execute and cause changes to memory. In so doing it can allow or disallow other productions which are looking for their conditions to be satisfied. Both the spread of activation and the selection of productions are parallel processes whose rates are controlled by "strengths" of network links and individual productions. An important aspect of this parallelism is that it is possible for multiple productions to be applied in a cycle through the set of productions. Much of the early work on the ACT system was focused on developing computational devices to reflect the operation of parallel, strength-controlled processes and working out the logic for creating functioning systems in such a computational medium.

We have successfully implemented a number of small-scale systems that model various psychological tasks in the domain of memory, language processing, and inferential reasoning. A larger scale effort is underway to model the language processing mechanisms of a young child. This includes implementation of a production system to analyze linguistic input, make inferences, ask and answer questions, etc. Also a great deal of effort is being given to developing learning mechanisms that will acquire and organize the productions for this language processing. This learning program attempts to acquire procedures from examples of the computations desired of the procedures. For instance, the program learns to comprehend and generate sentences by being given sentences and picture representations of the meaning of the sentences (actually hand encodings of the pictures). Although this effort is focused on induction of linguistic procedures, the hope is to develop a general model of induction of cognitive procedures and not to place any language-specificity into the induction procedures.

At the time of this report, we have completed the F version of ACT which is the system with learning capabilities. We are currently testing and tuning the system on a number of linguistic examples. Other projects which are progressing in earlier versions of ACT include use of spreading activation to model semantic disambiguation, modeling of the reading process, and modeling of solutions to word arithmetic problems.

D. Current list of project publications:

- [1] Anderson, J.R. Computer simulation of a language acquisition system: A second report. In D. LaBerge and S.J. Samuels (Eds.). Perception and Comprehension. Hillsdale, N.J.: L. Erlbaum Assoc., 1976.
- [2] Anderson, J.R. Language, Memory, and Thought. Hillsdale, N.J.: L. Erlbaum, Assoc., 1976.
- [3] Anderson, J.R. Induction of augmented transition networks. Cognitive Science, 1977, in press.
- [4] Anderson, J.R. & Kline, P. Design of a production system. Paper to be presented at the Workshop on Pattern-Directed Inference Systems, Hawaii, May 23-27, 1977.
- [5] Anderson, J.R., Kline, P. & Lewis, C. Language processing by production systems. To appear in P. Carpenter and M. Just (Eds.). Cognitive Processes in Comprehension. L. Erlbaum Assoc., 1977.
- [6] Kline, P.J. & Anderson, J.P. The ACTE User's Manual, 1976.

II. Interaction With the SUMEX-AIM Resource

The SUMEX-AIM resource is superbly suited for the needs of our project. We have made the most extensive use of the INTERLISP facilities and the facilities for communication on the ARPANET. We have found the SUMEX personnel extremely helpful both in terms of responding to our immediate emergencies and in providing advice helpful to the long-range progress of the project. Despite the fact that we are on the other side of the continent, we have felt almost no degradation in our ability to do research. We find we can easily list on the terminal a small portion of programs under modification. The willingness of SUMEX mail listing has also meant we can keep relatively up-to-date records of all programs under development.

A unique east coast advantage of working with SUMEX is the low loading of the system during the mornings. We have been able to get a great deal of work done during these hours and try to save our computer-intensive work for these hours.

We have found our one AIM work shop so far (1976) a very useful opportunity to meet with colleagues and exchange ideas.

A particularly striking example of the utility of the SUMEX resource was illustrated in the move from Michigan. In the summer of 1976 Anderson moved to Yale and Greeno to Pittsburgh. There was no loss at all associated with having to transfer programs from one system to another. At Yale we were programming the day after we arrived. The SUMEX link has also permitted continued collaboration with Greeno.

6.2.2 CHEMICAL SYNTHESIS PROJECT (SECS)

SECS - Simulation and Evaluation of Chemical Synthesis

W. Todd Wipke
Department of Chemistry
University of California at Santa Cruz

I. Summary of Research Program

A. Technical Goals.

The long range goal of this project is to develop the logical principles of molecular construction and to use these in developing practical computer programs to assist investigators in designing stereospecific syntheses of complex bio-organic molecules. Our specific goals this past year focused on improvement of the library of chemical transforms, completion of the perception of molecular symmetry and integrating the use of symmetry information throughout SECS including the strategy module. We also wanted to improve the execution speed of SECS, and the speed of graphical interaction over remote communication lines. We planned to simplify the program from the user's viewpoint by including automatic file failsafing, improvement of HELP commands, and non-fatal handling of all errors, as well as production of user's manuals for operation of the program and the writing of chemical transforms. Additionally we intended to initiate applications of SECS to the areas of biosynthesis and metabolism of compounds, as well as phosphorus chemistry. Finally we hoped to improve the strategic constraints and controls that guide SECS in growing a synthesis tree.

B. Medical Relevance and Collaboration.

The development of new drugs and the study of how drug structure is related to biological activity depends upon the chemist's ability to synthesize new molecules as well as his ability to modify existing structures, e.g., incorporating isotopic labels into biomolecular substrates. The Simulation and Evaluation of Chemical Synthesis (SECS) project aims at assisting the chemist in designing stereospecific syntheses of biologically important molecules. The advantages of this computer approach over a manual approaches are manifold: 1) greater speed in designing a synthesis; 2) freedom from bias of past experience and past solutions; 3) thorough consideration of all possible syntheses using a more extensive library of chemical reactions than any individual person can remember; 4) greater capability of the computer to deal with the many structures which result; and 6) capability of computer to see molecules in graph theoretical sense, free from bias of 2-D projection.

SECS was designed to be able to apply any kind of chemical transformation, and because of this generality we see SECS finding application in biogenesis and metabolism (see section II A below). The objective of using SECS in biogenesis is to predict possible biogenetic pathways for a given natural product and also

to predict related compounds which might also co-occur in nature. This can be a great aid in searching for new natural products and in structure elucidation.

The objective of using SECS in metabolism is to predict the plausible metabolites of a given xenobiotic in order that they may be analyzed for possible carcinogenicity. Metabolism research may also find this useful in the identification of metabolites in that it suggests what to look for, and in the identification of possible metabolic pathways connecting a metabolite to a xenobiotic.

C. Progress and Accomplishments.

RESEARCH ENVIRONMENT: At the University of California, Santa Cruz, we have a GT-40 graphics terminal connected to the SUMEX-AIM resource by a 1200 baud leased line and a TI 725 thermal printing teletype connected via TYMNET at 300 baud. UCSC has only a small IBM 370/145 and a PDP-11/45 (limit of 12 K words per user) available, both of which are unsuitable for this research. From July until December our research group had to occupy temporary space during renovation, but is now finally in permanent space in Thimann Laboratories where we have close collaboration with other organic chemists.

CHEMICAL TRANSFORMS: The library of chemical transforms has been reorganized and reevaluated during the past year by Mr. Dolata, a student of Professor D.A. Evans of Cal Tech. New reactions were added and the scope and limitations of others were updated and leading references provided. Additionally, Merck, Sharp, and Dohme Research Laboratories provided revisions of many transforms which a group of 25 synthetic chemists had carefully researched.

SYMMETRY: An efficient algorithm for recognizing molecular symmetry was developed last year. This year that algorithm has been tested against all possible molecular point groups and a few problems which developed were corrected. The algorithm has been documented and initial studies begun on actually determining the point group of a molecule. The symmetry group is now utilized in conjunction with the symmetry of a chemical transform so the transform is applied in all possible unique ways, to generate a non-redundant set of precursors. This symmetry of course takes into account stereochemistry of saturated centers and double bonds. We have surveyed literature syntheses for examples of existing heuristics based on symmetry which can be used for automatically generating high level strategies. This information has never been pulled together before and should make an interesting contribution also to organic synthesis.

STRATEGIC CONTROL: Last year we began developing an implementation of strategic control for SECS, and a simple language for expressing strategies independent of chemical transforms. Since these strategies contain expressions which refer to the molecular structure, it was also necessary to incorporate symmetry here too. For example, if a particular bond is designated as strategic to break, but a transform breaks another bond, the strategy is still satisfied if the two bonds are equivalent by symmetry. This problem becomes more complex when pairs of bonds are specified and when there are logical connectives (AND, OR, XOR, and NOT) involved. This has however been solved. Other changes since last year include a completely new user interface to strategy to allow error

correction and very easy modification of goals. Finally quantitative experiments have been performed to measure the effect of developing a synthesis tree with various types of strategic constraints. The net result of this work is that the user can more easily constrain SECS now to work only in areas which the user decides are worthwhile, consequently fewer precursors are generated which the user would delete.

USER INTERFACES: Users of SECS had difficulty understanding how to copy files into work areas in order to save or restore synthesis trees. Now SECS does all file manipulation, eliminating the problem. Further SECS now automatically failsafes the synthesis tree at key points so that in the event of machine or communication failure the user can automatically restart his analysis from the last key point. Considerable modifications were made to the graphical interface for increasing readability and speed of interaction. Over long slow communication lines (which happens to be the way most SECS users are accessing the program) interactive graphics must be done with care, minimizing the amount and frequency of picture transmission, in order to achieve even tolerable man-machine communication. Lastly, we have implemented appropriate input procedures to eliminate the possibility of a fatal crash from user input errors. According to user reports this was a major problem.

PHOSPHORUS CHEMISTRY: Graphical input and output procedures were developed for entering the stereochemical configuration of a trigonal bipyramid (TBP) phosphorus atom and for producing a correct structural diagram from the machine's internal representation. The SEMA algorithm for generating a stereochemically unique name was extended to deal with the 20 possible configurations for each TBP center, including the ability to recognize enantiomers. The ALCHEM language for representing chemical transforms was extended to facilitate manipulation of TBP's, including changes from trigonal and tetrahedral configurations to square base pyramid and TBP. Queries may deal with apicophilicity, and axial or equatorial orientation. The fine details of phosphorus chemistry such as the fact that groups entering or leaving the phosphorus coordination sphere normally do so from the apical position. Pseudo rotation, apicophilicity, and strain energy are considered in evaluating the stable TBP configurations and in checking for identical structures. A library of phosphorus chemistry is now being prepared in collaboration with a group at the University of Strasbourg, France.

COMPUTER-AIDED ELUCIDATION OF BIOGENETIC PATHWAYS: Although a great amount of effort has been spent on various areas of biogenesis, there have been few attempts to develop general techniques for the elucidation of biogenetic schemes. As a result, the formulation of biogenetic schemes has often been criticized for its lack of rigor and explicit criteria. Our approach is to develop general techniques which lead to the postulation of plausible biogenetic pathways, using the SECS as an aide in obtaining and analyzing solutions to this complex problem. It is our hope this application of computer problem solving techniques will not only uncover new ways of recognizing and evaluating biogenetic pathways but also provide added support to deductions made from biogenetic schemes, such as the generality of a scheme which may be tested in only a few species.

With the proper input information and goals well defined there may be explicit rules to guide the chemist to plausible biogenetic pathways for a particular natural product. Unfortunately, the vast majority of solutions to this problem are determined by a combination of the experienced natural products

chemist's ability to consider the most important rules involved and his unique set of experience-based prejudices. There may be some means to represent and utilize all of the known relevant rules, data and possibly even experience-based prejudices to arrive at the best plausible pathways. The most precise method for representing, developing and testing such a theory is in the form of a computer program. To implement such a computer program, known rules and constraints must be clearly defined, then those that are applicable can be applied at each step of the analysis toward the desired goal. This will keep the solution pathways logically pure and insure that all alternatives which satisfy the rules and constraints are considered. This guarantee of completeness simply can not be made using hand analysis.

A new reaction library containing biogenetic transformations have been written. After inputting a natural product the program will apply the biogenetic transforms which fit the natural product. This generates a set of plausible biogenetic precursors to the target natural product. By continuing this process with the precursors generated, the plausible biogenetic pathways for the natural product can be discovered.

The structures of marine natural products were entered into the program and the plausible biogenetic pathways for these compounds were generated and analysed. Biogenetic pathways which had been proposed in the literature were among the pathways discovered, as were other plausible pathways which would now have to be considered. The success we attained in this research effort verified the applicability of the SECS program as an aid in the analysis of metabolic pathways.

COMPUTER-AIDED PREDICTION OF METABOLITES FOR CARCINOGENICITY STUDIES: We have initiated a research project in collaboration with the Chemical Carcinogenesis group at the National Cancer Institute. The objective of this research is to establish a computer program by which a biochemist or metabolism expert can explore the metabolism of a chemical compound. The investigator enters the substrate molecule by interacting with an input and structure editing module.

Then the program will apply the biological transforms which "fit" the structure, taking into consideration all the context information (2-D, 3-D, and electronic) available about the transform and all perceived information about the structure. This will generate a set of metabolites which are one step away from the substrate structure. The metabolites will be ranked according to expected probability or yield. The exact parameters which should be monitored will be determined during the course of this research. An evaluation module may then screen these metabolites according to criteria specified by the investigator. Duplicate metabolites arising from different pathways will be labelled to indicate that fact. Finally the investigator will be shown the set of metabolites together with data about the transform which produced each one and the values of the parameters being monitored.

The investigator may select one metabolite for further metabolism or may request that all be processed for a specified number of steps. In this way a "tree" of metabolites is produced and displayed. The entire state of the user's tree may be saved to permit continuation of the analysis at another time. Exploration of the metabolism tree will be predominately guided interactively by

the expert investigator. We feel that at this stage of development of the field of metabolism and carcinogenicity that interactive guidance by the expert is necessary. There are many areas where the theory is very thin and a given biological transformation may have been observed for only a few substrates. When this transform is applied to a new substrate, some unrealistic metabolites may be generated owing to the deficiency of contextual information and constraints. An expert is necessary to prune the tree and prevent the automatic processing of those unreasonable intermediates. It is much more efficient for the expert to do this pruning as the tree is being grown, rather than later after an enormous tree has been completed.

At some point either during tree generation or at the end, the metabolites will be passed to another program which will identify those metabolites which are identical or "similar" to known carcinogens. Those will be so marked in the tree.

Presently, the major task is the acquisition of the metabolism knowledge base, i.e. the writing of the transformation library to be utilized. Metabolism experts at the National Institute of Health are gleaning this information from both their own research and the metabolism literature. This information will be encoded and the first testing of this new application for the SECS program will begin in June 1977.

D. Current List of Project Publications

- W.T. Wipke and P. Gund, "Simulation and Evaluation of Chemical Synthesis. Congestion: A Conformation Dependent Function of Steric Environment at a Reaction Center. Application with Torsional Terms to Stereoselectivity of Nucleophilic Additions to Ketones," J. Am. Chem. Soc., 98, 8107(1976).
- W.T. Wipke, G. Smith and H. Braun, "SECS-Simulation and Evaluation of Chemical Syntheses: Strategy and Planning," ACS Symposium Proceedings, 1977.
- W.T. Wipke, Computer Planning of Research in Organic Chemistry, Proceedings of the Third International Symposium on Computers in Chemical Education, Research, and Technology, Caracas, Venezuela, 1976.

S.A. Godleski, P.v.R Schleyer, E. Osawa, and W.T. Wipke, "The Systematic Prediction of the Most Stable Neutral Hydrocarbon Isomer," J. Am. Chem. Soc., 99, 0000(1977).

F. Choplin, R. Marc, G. Kaufmann, and W.T. Wipke, "Computer Design of Synthesis in Phosphorus Chemistry. Automatic Treatment of Stereochemistry," J. Am. Chem. Soc., 99, 0000(1977).

Manuals:

SECS Users Manual, June 1976.
SECS Users Guide, Aug 24, 1976.
ALCHEM Tutorial, Sep 21, 1976.

II. Interactions with SUMEX-AIM Resource

A. Examples of Collaborations and Medical use of Programs via SUMEX.


SECS is available in the GUEST area of SUMEX and has been accessed experimentally by many others as well. Professor R. V. Stevens (UCLA) explored some syntheses of lycapodine while visiting Santa Cruz and as a result has requested UCLA to obtain a graphics terminal so he and others at UCLA can access SECS via SUMEX. Professor W. G. Dauben's group (Berkeley) has utilized the SECS model builder on SUMEX is now extending the capabilities of that module of SECS. Mr. Mel Spann of the National Library of Medicine toxicology program is collaborating with us in developing a metabolism library for the metabolism of catechol amines. Also collaborating with us on metabolism are Drs. Ted Gram from Guarino's lab, Harry Gelboin, Dhiren Thakken and Harukiko Hagi from Jerina's lab, Lance Pohl from Gillette's lab, Sidney Nelson from Mitchell's lab, Lionel Poirier from Weisburger's lab, and Ken Chu and Sidney Siegel all of whom are from the National Cancer Institute. Dr. Steve Heller of the EPA and Dr. G.A. Milne of the National Heart and Lung Institute have expressed interest in putting SECS on the Cyphernetics network as a part of the NIH chemical information system. Restrictions on the allowed core image on that system have so far held up the negotiations.

For the past two years SECS has been available over TELENET from First Data Corporation and has been accessed by industry: Squibb; Merck, Sharp and Dohme; Pfizer; Searle; Lederle Labs; FMC; and recently 3M Corporation and Stauffer. Dr. Beryl Dominy of Fizer recently presented a paper before the Pharmaceutical Manufacturer's Association entitled "SECS and the Information Scientist" in which he describes his experiences with SECS, including an example where a synthetic chemist was having difficulty with a particular synthesis, he then went to SECS for possible solutions. SECS suggested another route as being better and indeed that is what he found when he tried it later in the lab.

The availability of SECS on SUMEX-AIM has also served health-related research at the University of California, Santa Cruz. Model building using the SECS model builder is being performed for Professor Edward Dratz (UCSC) to generate conformations of fatty acids isolated from visual membranes ("Structure

and Function of Visual photoreceptors," EI00175), and for Professor Howard Wang (UCSC) to study how conformations of steroids may affect the local anesthetic - membrane interaction ("Role of Membrane Proteins in Local Anesthetic Action," GM22242).

We have assisted Professor J. E. McMurry in his synthetic work towards Aphidicholine and Digitoxigenin by using the model builder for predicting possible reaction pathways. An example is given below, where the conformation of the epoxy-ylide was calculated along with the strain energies of the two possible closure products.



Utilizing the SECS model builder, we have shown that attack on the epoxide to form the fused system should be much more favorable than attack to form the bicyclo compound. Similar studies have been undertaken to predict the stereochemistry resulting from the acid catalyzed cyclization of McMurry's Digitoxigenin precursor (HL-18118 "Total Synthesis of Cardiac Aglycones."): application of SECS using a special library of cationic sigmatropic rearrangement transforms generated the possible products which facilitated identification of some of the side products in the early cyclization experiments. We have also collaborated in the biogenesis work with Professor Phil Crews (UCSC) in marine natural product biogenesis. Dr. Wipke has also used several SUMEX programs such as CONGEN in his course on Computers and Information Processing in Chemistry.

B. Examples of Sharing, Contacts and Cross-fertilization with other SUMEX-AIM projects.

In collaboration with Dr. Ray Carhart and Dr. Dennis Smith of the DENDRAL/CONGEN Project, a Computers in Chemistry Workshop was held at U.C. Santa Cruz on the weekend prior to the Fall 1976 American Chemical Society National Meeting held in San Francisco. The workshop attracted participants representing all parts of the chemical community, academia, industry and government. Morning lecture/discussion sessions introduced the SECS and CONGEN programs running on

SUMEX and the afternoon and evening sessions allowed "hands-on" experience for the participants. The response of the workshop participants was a very positive one with many participants showing so much interest that future collaboration and/or use of the powerful non-numerical computing tools available on SUMEX was discussed.

The SECS project has held joint research group meetings at Stanford with the DENDRAL and AI groups to discuss common problems and research goals. This has been very rewarding since the groups are complementary in orientation. These joint meetings also let the members meet in person after having met on-line on the network. Last year's AIM Conference at Rutgers was also a valuable experience, which allowed us to meet people interested in similar problems in different disciplines. It was particularly useful to have the opportunity to talk with experts designing new languages for knowledge representation and to hear them compare their systems.

6.2.3 HIGHER MENTAL FUNCTIONS PROJECT

Modeling of Higher Mental Functions

Kenneth M. Colby, M.D.
Professor of Psychiatry and Biobehavioral Sciences
University of California at Los Angeles

I) Summary of Research Program

A. Technical Goals:

There are three technical goals of the Higher Mental Functions Project:

- (1) To improve and "therapeutically" experiment with a computer simulation of paranoid processes in order to make treatment recommendations to clinicians based on experience with the model.
- (2) To develop a new taxonomy of psychiatric patients based on the conceptual patterns appearing in accounts of their illnesses.
- (3) To develop an intelligent speech prosthesis for patients suffering from communication disorders.

B. Medical Relevance and Collaboration:

The Higher Mental Functions Project is located in the Neuropsychiatric Institute at UCLA. The medical relevance of its research concerns the fields of psychiatry and neurology. The Project collaborates with clinicians and investigators in psychiatry, neurology, the neural sciences and neurolinguistics.

C. Progress Summary:

We have improved the paranoid model to the point where it can be utilized for therapy experiments. (The model has now passed a true Turing Test in which it cannot be distinguished from real patients.)

The taxonomy effort is just under way, using the language recognition program which serves as the front end of the paranoid model. This program will have to be added to and modified to serve the purpose of finding and classifying the conceptual patterns appearing in patients' accounts of their illnesses.

We have interfaced a micro-processor with a voice-synthesizer to provide a speech prosthesis for patients unable to speak. The next step is to write an "intelligent" algorithm which attempts to figure out what the patient is trying to say from his partial input information.

II. Interactions with the SUMEX-AIM Resource

A. Collaborations:

The project collaborated with Professor Jon Heiser, Department of Psychiatry, University of California, Irvine, and consulted with Professor Robert K. Lindsay, Department of Psychology, University of Michigan, in conducting a Turing Test of the paranoid model. Other users of SUMEX have received advice and suggestions regarding their problems as well as opportunities to contrast their simulations with ours. We have benefitted greatly from others' comments on the adequacy and inadequacy of our paranoid model.

B. Sharing, etc.:

Members of the project have participated in two workshops held at Rutgers, presenting several papers, chairing panels, and conducting discussion groups. Informal discussions with large numbers of workers in Artificial Intelligence in Medicine have led to a helpful sharing of ideas and techniques. SUMEX is valuable to us as a communication channel combining the advantages of a telephone and the U.S. mail without the disadvantages of either. For widely scattered researchers, it facilitates the intimate, low-level communication which is normally accomplished in hallways or around water coolers. The individual discussions are not very profound, but the cumulative effect subtly improves our research.

The existence of SUMEX as an independent project naturally relieves numerous researchers of the burden of separately financing and staffing a large computer facility.

D. Up-to-date List of Publications:

- Colby, K.M., Parkison, R.C. and Faught, B. Pattern-matching Rules for the Recognition of Natural Language Dialogue Expressions. Am. J. Computational Linguistics, Microfiche 5, Sept., 1974.
- Colby, K.M. Clinical Implications of A Simulation Model of Paranoid Processes. Archives of General Psychiatry, 33, 854-857, 1976.
- Faught, W., Colby, K.M. and Parkison, R.C. Inferences, Affects and Intentions in A Model of Paranoia. Cognitive Psychology, 9, 153-187, 1977.
- Colby, K.M. An Appraisal of Four Psychological Theories of Paranoid Phenomena. J. of Abnormal Psychology, 86, 54-59, 1977.
- Parkison, R.C., Colby, K.M. and Faught, W.S. Conversational Language Comprehension Using Integrated Pattern Matching and Parsing. Artificial Intelligence (In Press) 1977.
- Colby, K.M., Christinaz, D. and Graham, S. A Computer-driven, Personal, Portable and Intelligent Speech Prosthesis for Aphasic Disorders. Brain and Language (In Press) 1977.
- Colby, K.M. On the Way People and Models Do It. Perspectives in Biology and Medicine (In Press) 1977.

Heiser, J., Colby, K.M., Faught, W. and Parkison, R.C. Testing Turing Test (Forthcoming).

Faught, W.S. Conversational Action Patterns in Dialogs. Proceedings of the Workshop on Pattern-directed Inference Systems, May, 1977.

6.2.4 INTERNIST PROJECT

INTERNIST - Diagnostic Logic Project

J. Myers, M.D. and H. Pople, Ph.D.
University of Pittsburgh

I. SUMMARY OF RESEARCH PROGRAM

A. Objectives

The principal objective of this research project has been and continues to be the development, evaluation, and implementation of a computer-based diagnostic consultation system for internal medicine. This work, which was initiated at the University of Pittsburgh approximately six years ago, has been supported for the past three years by a grant from the Bureau of Health Resources Development. A heuristic diagnostic program called INTERNIST has been developed, along with an extensive medical database now comprising more than four hundred disease categories and two thousand manifestations of disease. The system has been tested with a wide variety of difficult clinical problems: cases published in the medical journals, CPC's, and other interesting and unusual problems arising in the local teaching hospitals. In the great majority of these test cases, the heuristic INTERNIST program has proved to be effective in sorting out the pieces of the puzzle and coming to a correct diagnosis. In some cases, as many as six distinct disease entities have been identified correctly.

We believe that by the time of the expiration of the BHRD grant in June, 1977, our original objective, which was to develop a system providing expert diagnostic capability with regard to the major diseases of internal medicine, will have been accomplished to the extent possible in the current laboratory framework.

At that time, we propose to initiate a broader collaboration, which will invite the participation of remote users in

- (a) further evaluation of the INTERNIST programs and data-base.
- (b) development of specialized data-bases and procedures for various medical subspecialties.
- (c) refinement of the user interface.
- (d) investigation of alternate uses of the INTERNIST data-base.

We believe that the expansion of the experience base of INTERNIST users, which will result from this type of collaboration, will significantly enhance the further course of INTERNIST development.

B. Progress Summary

Expansion of the medical data-base to encompass new areas of disease is an on-going activity of the project. Much of this work is carried out by medical students who elect to take part in the project as part of their fourth year clinical rotation, with the period of participation varying from 6 to 18 weeks.

Each student is assigned a group of diseases, usually in a specific clinical area, for study. The literature on a disease is studied exhaustively for all quantitative data available. Frequently clinical experts on the faculty are consulted, particularly about controversial data. The student compiles a complex list of the manifestations of the disease under study and assigns tentative measures of strength of association.

The clinical principal investigator together with any other clinicians working on the project then review the data exhaustively in order to assure the appropriateness and completeness of the disease profile.

The profile is then entered into the computer and tested for completeness and reliability against a typical or "textbook" example of clinical cases. If available, other cases of the disease from the floors of our university hospital and from published cases such as the clinical-pathological conferences from the New England Journal of Medicine and the American Journal of Medicine are also used. Further refinement occurs in the course of the continued use of the data-base.

In addition to this data-base development, work on a refined diagnostic program has also been an on-going activity during this period.

The present INTERNIST process employs a 'problem - formation' heuristic, which identifies one of perhaps several problems in a clinical case as its initial focus of problem-solving attention. Although only one problem is considered at a time, the process recycles after each problem is solved, thereby uncovering the entire complex of diseases present. In the great majority of clinical cases tested, this strategy of iterative problem formation and solution has proved to be effective in sorting out the complexities of a case and rendering a correct diagnosis. In many respects, however, it seems clear that performance could be significantly enhanced if the program were to attend to the various component problems and their inter-relationships simultaneously. Use of a more global problem - formation strategy could be expected to yield more rapid convergence on the correct diagnosis in many cases, and in at least some cases to prevent missed diagnoses.

Alternative problem formation strategies that exploit the type of pseudoparallel processing facilitated by the INTERLISP 'spaghetti stack' are presently being investigated. We believe that this research will also set the stage for subsequent development of a therapeutic management component of the INTERNIST consultation facility; however at the present time it is not possible to project a precise timetable for the development of these additional capabilities.

C. Publications

1. Pople, H.E., Myers, J.D., & Miller, R.A., "The DIALOG Model of Diagnostic Logic and its use in Internal Medicine". Proceedings of the Fourth International Joint Conference on Artificial Intelligence, Tbilisi, USSR, September 1975.
2. Pople, H.E., "Artificial-Intelligence Approaches to Computer-based Medical Consultation, Proceeding IEEE Intercon, New York, 1975.
3. Pople, H.E., "The Syntheses of Composite Hypotheses in Diagnostic Problem Solving: An Exercise in Hypothetical Reasoning". Proceedings of the Fifth International Joint Conference on Artificial Intelligence, August 1977 (forthcoming).

II. INTERACTION WITH SUMEX-AIM RESOURCE

A. Medical Use of Programs and Collaborations

Because of the research and development nature of our work on the INTERNIST system over the past several years, we have been somewhat limited in our ability to establish wide-spread collaborations. However, members of the medical house staff in the local hospitals having some prior experience with the project have continued to work with INTERNIST while pursuing their medical training. In addition, project staff often have occasion for interaction with individuals and groups who have interest in the characteristics of the diagnostic system from both medical and computer science perspectives. Future plans for more extensive collaboration are discussed in section III.

B. AIM Interactions

We have benefitted considerably from interactions with other members of the SUMEX-AIM community. In June '76 we participated in the AIM workshop at Rutgers, which provided an excellent perspective as to what else is going on in the field. During the past several months we have had useful exchanges with Randy Davis, Victor Yu, and John Foy, three individuals participating in the MYCIN project. In addition, we rather routinely interact with SUMEX staff regarding fine points and problems relating to our use of system facilities.

The opportunity to keep abreast of developments in a fast changing field is one of the principal benefits to be derived from the collegial environment fostered by SUMEX-AIM.

6.2.5 MEDICAL INFORMATION SYSTEMS LABORATORY

MISL - Medical Information Systems Laboratory

M. Goldberg, M.D. and B. McCormick, Ph.D.
University of Illinois at Chicago Circle

I) SUMMARY OF RESEARCH PROGRAM

A.) TECHNICAL GOALS

The Medical Information Systems Laboratory (MISL) was established under grant HM-0114 in Chicago to pursue three activities: i) Construction of a database in ophthalmology, ii) Clinical knowledge system support, and iii) Network-compatible database design. Priorities in year 04 of MISL's operation are the same as in previous years: investigations into how to construct a database in ophthalmology, and into distributed database design, are ancillary to the exploration of a clinical knowledge system to support clinical decision making. We are developing ways to get reliable clinical information into the ophthalmic database primarily because we are interested in getting out significant clinical decision support.

B) APPROACH AND MEDICAL RELEVANCE

B.1) Construction of the database in Ophthalmology

A specific aim of this project is to construct a workable database in ophthalmology, using the outpatient population of the Illinois Eye and Ear Infirmary. We view this database as a testbed for developing clinical decision support systems. The Ophthalmology Department of the Illinois Eye and Ear Infirmary provides an excellent environment for evaluating new techniques for capturing and using clinical information.

B.2) Clinical knowledge support system

The goals for clinical knowledge system development are to provide a flexible user interface for a prototype relational database system, to devise means of accessing alphanumeric and pictorial information stored in the database system, and to provide efficient means for logically restructuring a database so that it can be adapted to different operating environments in a network-compatible distributed medical information network.

No clinical database, however, has intrinsic significance beyond its ability to support the diagnosis and management of disease. Additional goals for the clinical knowledge system are therefore to devise computer-based consultation systems for glaucoma and selected retinal/choroidal diseases, and to provide

formal models which permit the relational development and evaluation of rule-based consultation systems containing 2,000 - 10,000 rules. In recognition that a continuum exists between physician-guided decision support and computer-based consultation, we choose to describe these services as a Clinical Knowledge System: a consortium of a clinical database and rules for its interpretation.

C) PROGRESS SUMMARY (INCLUDING ITEMS OF INTEREST TO SUMEX-AIM COMMUNITY ONLY)

C.1) The database in ophthalmology

Physician terminals and interfaces to ophthalmic instruments have been positioned in the general eye clinic and several key ophthalmic subspecialty clinics. Systematic, modular hardware and software for clinical source data acquisition have been established. The clinical support system computer will shortly be transferred to the newly dedicated Goldberg Research Center, adjacent to the Illinois Eye and Ear Infirmary. We look forward to stabilizing the hardware configuration, telecommunication linkages and software support.

C.2) Clinical knowledge system support

C.2.a) Development of the relational database includes the following:

- A user interface through which unsophisticated users communicate with the database.
- An intelligent coupler that serves as an intermediary between the end user and the distributed database system. The coupler listens to the user's retrieval requests; helps the user formulate his requests correctly; efficiently translates user's retrieval requests into a network-compatible retrieval command language; and obtains authorization from the system for data retrieval and/or update.
- Tools for picture data management. Graphical indexing techniques are provided so that the clinical researcher and physician can easily retrieve pictorial/graphical information from the medical database.
- Means for logical database synthesis. This involves conversion of the user's view of the database into a logically coherent physical organization.

C.2.b) Development of a computer-based consultation system for diagnosis and management of glaucoma.

This involves on-going collaboration between Dr. Jacob Wilensky at MISL, and, through SUMEX-AIM, other investigators around the United States. Included are the original investigators in glaucoma consultation: Dr. Casimir Kulikowski (Rutgers), Dr. Shalom Weiss (Mt. Sinai Hospital, NY), and Dr. Aaron Safir (Mt. Sinai Hospital).

C.2.c) Development of a consultation system for diagnosis and management of retinal/choroidal diseases.

A design has been proposed (in Walser and McCormick, see below) for MEDICO, a consultation system that advises non-expert physicians in the management of chorioretinal diseases. In addition, a major subsystem of MEDICO, responsible for mediating the acquisition and organization of rules, has been implemented.

C.2.d) Formal models for consultation systems.

Petri nets have been studied, primarily by Murata (see below), as a formal representation for interacting parallel processes. Petri nets are similar to causal networks, as described by Kulikowski and Weiss at Rutgers, except that, with Petri nets, cyclic activity is easily represented. The similarity between Petri nets and inference nets has also been noted (Walser and McCormick). The utility of the Petri net framework for modelling physical processes was explored by Walser, with the construction of a simulated coffee maker. Further studies are planned.

D.) LIST OF MISL PUBLICATIONS

- Chang S. K., Donato N., McCormick B. H., Reuss J., and Rocchetti R. (1977) A relational database system for pictures. Proc. IEEE Workshop on Picture Data Description and Management, April 20-22, 1977, Chicago, Illinois.
- Chang S. K. and Cheng W. H. (1976) A database skeleton and its application to logical database synthesis. MISL report M.D.C. 1.1.17.
- Chang S. K. and McCormick B. H. (1975) An intelligent coupler for distributed database systems. MISL report M.D.C. 1.1.7.
- Malone, J. E. (1976) Interval generalization of structure representation. MISL report M.D.C. 1.1.22.
- Malone J. E. (1975) User's guide to uniclass cover synthesis. MISL report M.D.C. 4.4.1.
- Malone J. E. (1975) Addendum to AQVAL/1 (AQ7), part 1: User's guide and program description. MISL report M.D.C. 4.4.1.
- Manacher G. K. (1977) The case for strong loops and selection structures in ordinary computer languages. MISL report M.D.C. 1.1.21.
- Manacher G. K. (1975) On the feasibility of implementing a large relational data base with optimal performance on a minicomputer. Proc. International Conference on Very Large Data Bases, Framingham, Mass.
- McCormick B. H. and Nordmann B. J. Jr. (1977) Modular asynchronous control design. Forthcoming in IEEE Transactions on Computers. Also MISL report M.D.C. 1.1.25.

- McCormick B. H. and Amendola R. C. (1977) Cytospectrometers for subcellular particles and macromolecules: design considerations. Presented at Workshop on Theory, Design and Biomedical Applications of Solid State Chemical Sensors, Case Western Reserve University, March 23-30, 1977. Also MISL report M.D.C. 1.1.24.
- McCormick B. H. and Wilensky J. (1975) Clinical knowledge acquisition: design of a relational data base in ophthalmology. Proc. Second Annual Medical Information Systems Conference, Urbana, Ill.
- McCormick B. H., Goldberg M. F., and Read J. S. (1974) Clinical decision-making: design of a data base in ophthalmology. Proc. First Annual Medical Information Systems Conference, Urbana, Ill.
- Michalski R. S. and Chang S. K. (1976) A self-model for a relational database. MISL report M.D.C. 1.1.16.
- Michalski R. S. (1975) On the selection of representative samples from large relational tables for inductive inference. MISL report M.D.C. 1.1.9.
- Murata T. (1976) On liveness and other properties of E-Nets. MISL report M.D.C. 1.1.15.
- Murata T. (1976) Bibliography on Petri nets and related topics. MISL report M.D.C. 1.1.20.
- Murata T. (1976) A method for synthesizing marked graphs from given markings. Presented at 17th Annual Symposium on Foundations of Computer Science, October 25-27, Houston, Texas.
- Murata T. (1976) On deadlock and the liveness of E-nets. Presented at the 17th Annual Symposium on Foundations of Computer Science, October 25-27, Houston, Texas.
- Murata T. (1975) State equation, controllability, and maximal matchings of Petri nets. MISL report M.D.C. 1.1.10.
- Murata T. and Church R. W. (1975) Analysis of marked graphs and Petri nets by matrix equations. MISL report M.D.C. 1.1.8.
- Vere S. A. (1975) Induction of concepts in the predicate calculus. Proc. Fourth IJCAI.
- Vere S. A. (1975) Relational production systems. Forthcoming in Artificial Intelligence. Also MISL report M.D.C. 1.1.5.
- Walser R. L. and McCormick B. H. (1976) Organization of clinical knowledge in MEDICO. Proc. Third Illinois Conference on Medical Information Systems, Urbana, Ill.
- Walser R. L. and McCormick B. H. (1977) A system for priming a clinical knowledge base. Forthcoming in Proc. 1977 National Computer Conference, June 13-16, Dallas, Texas.

II) INTERACTION WITH SUMEX-AIM RESOURCE

A.) COLLABORATION

Major collaboration at present is through the ONET, involving the ophthalmology departments of five medical schools. Dr. Jacob Wilensky is actively engaged in evaluating and modifying the Glaucoma Consultation Program, written originally by Shalom Weiss.

6.2.6 RUTGERS COMPUTERS IN BIOMEDICINE

Rutgers Research Resource - Computers in Biomedicine

Principal Investigator: Saul Amarel
Rutgers University, New Brunswick, New Jersey

I) SUMMARY OF RESEARCH PROGRAM

A) Goals and Approach

The fundamental objective of the Rutgers Resource is to develop a computer based framework for significant research in the biomedical sciences and for the application of research results to the solution of important problems in health care. The focal concept is to introduce advanced methods of computer science - particularly in artificial intelligence - into specific areas of biomedical inquiry. The computer is used as an integral part of the inquiry process, both for the development and organization of knowledge in a domain and for its utilization in problem solving and in processes of experimentation and theory formation.

The Resource community includes 48 researchers - 30 members, 8 associates and 10 collaborators. Members are mainly located at Rutgers. Collaborators are located in several distant sites and they interact, via SUMEX-AIM, with Resource members on a variety of projects, ranging from system design/improvement to clinical data gathering and system testing. At present, collaborators are located at the Mt. Sinai School of Medicine, N.Y.; Washington University School of Medicine, St. Louis, Mo.; Johns Hopkins Medical Center, Baltimore, Md.; Illinois Eye and Ear Infirmary, Chicago, Ill.; and the University of Miami.

Research in the Rutgers Resource is oriented to "discipline-oriented" projects in medicine and psychology, and to "core" projects in computer science, that are closely coupled with the "discipline-oriented" studies. Work in the Resource is organized in three AREAS OF STUDY; in each area there are several projects. The areas of study and the senior investigators in each of them are:

- (1) Medical Modeling and Decision Making (C. Kulikowski, A. Safir).
- (2) Modeling Belief Systems and Common-sense Reasoning (C.F. Schmidt, N.S. Sridharan).
- (3) Artificial Intelligence: Representations, Reasoning and System Development (S. Amarel)

In addition, the Rutgers Resource is sponsoring an Annual National AIM Workshop, whose main objective is to strengthen interactions between AIM activities, to disseminate research methodologies and results, and to stimulate collaborations and imaginative resource sharing within the framework of AIM. The second AIM Workshop was held near the New Brunswick Rutgers Campus on June 1-4, 1976. The third Workshop is scheduled for July 6-8, 1977.

B) Medical Relevance; Collaborations

A major part of our research is focusing on the development of computer based medical consultation systems. We are using artificial intelligence approaches in problems of: knowledge acquisition from experts in a medical specialty and from their clinical experience; the representation and management of these complex and changing data bases of medical knowledge within the computer; and the development of a sufficiently rich repertoire of reasoning strategies for diagnosis, prognosis, therapy selection, explanation and teaching. By linking such a system to a data base of prospectively chosen cases, we are in the position to provide a powerful tool for clinical research with built-in interpretative capabilities.

Our approach emphasizes the development and application of clinically useful models that describe the pathophysiology and dysfunction of diseases in a variety of tasks:

- a) Consultation embodying expert knowledge, which is expressed in terms acceptable to the clinician;
- b) Clinical research aid, assisting the investigator to;
 - i) Summarize and incorporate his knowledge, experience, and opinions into a computer system;
 - ii) Analyze his data, check it against that of other investigators, pooling it when appropriate to draw stronger conclusions based on the large sample of cases;
 - iii) Test, evaluate and modify the data base of models and decision strategies to produce an up-to-date summary of experience in his specialty.
- c) Screening and diagnosis, to aid nursing or paramedical personnel in performing routine decision procedures within restricted medical environments;
- d) Instruction to provide practitioners and support personnel with appropriate explanation and guidance in clinical decision-making.

A unique and novel aspect of our work is the creation of a network of clinical investigators to collaborate on the testing and continued development of the computer programs needed to accomplish the above tasks. During 1976, the ophthalmological network (ONET) of glaucoma investigators has grown and established itself, with several significant collaborative research projects currently underway. The consultation program for glaucoma using the causal associational network (CASNET) model developed within the Rutgers Resource, was jointly presented by the ONET members at the 1976 meeting of the Association for Research in Vision and Ophthalmology. An important new emphasis has been the incorporation into the consultation program of alternative expert opinions on subjects currently under debate. Dr. Douglas Anderson of the Bascom-Palmer Eye Institute at the University of Miami has joined ONET to provide such alternatives and strengthen the glaucoma model in certain important areas. The SUMEX-AIM shared computer resource has been essential to the activities of ONET.

The knowledge base and the strategies of our CASNET glaucoma consultation system are being strengthened and refined continuously in the ONET environment. The system is now at a point where it is considered by leading ophthalmologists as "highly competent to expert" in several subspecialties of glaucoma. The ONET group was confident enough about the system to demonstrate it at the October 1976 meeting of the American Academy of Ophthalmology and Otolaryngology. The reactions to the system were most favorable. The response of an independent sample of ophthalmologists taken at this meeting strongly emphasized the importance of the system for glaucoma research.

In addition to the main glaucoma research activities, the Resource has collaborated with the Mt. Sinai-Rutgers Health Care Computer Laboratory in the development of models for refraction and visual fields. These will be used by clinical prototype programs for guiding paramedical personnel in data acquisition and decision-making. These programs run on the PDP-11 computers of the clinical ophthalmological system at Mt. Sinai, which are to be linked to the PDP-10 at Rutgers for accessing the more complex models of disease when they are needed. The activities in conjunction with the Health Care Computer Laboratory reflect the more applied aspects of our work in the medical area.

The collaboration with Dr. R. Nordyke of the Straub Clinic on thyroid disease consultation systems has continued at a low level of activity during 1976.

In the area of Belief Systems, collaboration has continued with Professor Andrea Sedlak and her group at the University of North Carolina. This collaboration is focusing on developmental aspects of action perception.

In the AI Area we had extensive interactions with researchers in several institutions on problems of representation, problem solving systems, natural language processing, automatic programming, data base systems, and interactive systems. Contacts continued with the natural language group at BBN (Woods, Bruce) on the design of natural language processors for medical systems. Also, we had contacts with the Stanford-Xerox group (Winograd, Bobrow) which is involved in the development of KRL (Knowledge Representation Language).

Following the Rand Workshop on Biomedical Modeling (February 18-20, 1976), in which S. Amarel participated, preliminary contacts started with Dr. D. Garfinkel from the University of Pennsylvania in connection with possible applications of AI methods to the modeling of metabolic processes.

Our close contacts with the Stanford projects on Heuristic Programming (Drs. Buchanan, Feigenbaum, Lederberg) are continuing. The orientation and approach of these Stanford projects are very similar to ours. We continue to share with the investigators in DENDRAL and METADENDRAL a strong interest in computer-based methods of scientific inference and in AI ideas and techniques for representation of knowledge in computers, diagnostic problem solving and theory formation.

One of the significant collaborative developments this period was the joint work of Ed Feigenbaum and his students at Stanford, and Saul Amarel and his students at Rutgers, on the development of an AI Handbook. This handbook is being prepared on the SUMEX-AIM and RUTGERS-10 computers, and it is intended to

provide a network-accessible encyclopedic coverage of the AI field for the AIM community and AIM guests.

C) Progress Summary

1. Areas of Study and Projects

a) Medical Modeling and Decision-Making

The consolidation of the ophthalmological network (ONET) of collaborating glaucoma investigators using the SUMEX-AIM shared resource facility, the testing and improvement of the CASNET consultation system with the help of the collaborators, the design and implementation of a time-oriented database system and a set of analysis programs for aiding joint clinical research activities within ONET, and the development of a new knowledge-based consultation system (IRIS), represent the main achievements in the last year.

The network of investigators in glaucoma is designed to foster development of consultation systems that embody sufficient depth for knowledge and expert opinion in a variety of subareas to be useful as research and teaching tools. The collaborative activities, coordinated by Dr. A. Safir at Mt. Sinai, bring together selected scientist-users with complementary interests and strengths in different aspects of glaucoma, and Resource investigators who are concentrating on the development of new computer science methodologies in modeling and problem solving. During this period, there has been more extensive testing of the CASNET glaucoma consultation program. The collaborators had several meetings to discuss the structure of the glaucoma model and suggested many improvements and additions. A significant new capability of the program is the inclusion of alternative interpretations that capture differences of opinions among the experts on aspects of the model that are currently under debate.

A new development during this period has been the implementation of a time-sequenced data base for glaucoma, which has the dual purpose of aiding the clinical research of ONET collaborators and of providing a systematic means for evaluating and improving the performance of the consultation programs.

In the area of general methods and systems we have developed a multilevel-semantic network representation for characterizing disease processes, their anatomical descriptions and their taxonomic identification. This is used by a set of normative rules for diagnostic, prognostic and therapeutic reasoning, which results in a very general and flexible system for clinical consultation. A prototype model called IRIS is being developed using the glaucoma knowledge-base. We have also continued our investigations of other representation paradigms: a frame-based approach and the relationship to mathematical models of optics and refraction. Another subproject is concerned with developing methods of inference over network structures that will permit us to incorporate the results of clinical experience with different groupings of case-types into the models of consultation, aiding at the same time in the evaluation of the programs.

b) Modeling of Belief Systems and Common-Sense Reasoning

During this period a major achievement was the development and implementation of the AIMDS system. This is an MDS-based system that is specialized and augmented for use in modeling reasoning about actions. A noteworthy aspect of the system is the use of the MDS concepts of Consistency Conditions and Residues to guide frame instantiations and the drawing of further inferences from such frame instantiations.

The BELIEVER theory is a psychological model of the processes involved in the interpretation and common-sense reasoning about observed human actions. The AIMDS system is being constructed to provide a framework for formulating, studying and testing the BELIEVER theory. The computer system and the psychological theory are growing together, and they are strongly influencing each other's development. The domain of common-sense reasoning about actions represents a prototypical example of knowledge based reasoning. The richness of the psychological data that this theory must explain, namely, persons' linguistic descriptions and summarizations of everyday behavior, has forced us to think very carefully about how knowledge is to be represented and used. Out of this has emerged a general scheme that not only seems psychologically plausible but also appears to provide a useful framework for viewing a wide variety of problems of interpretation including medical diagnosis and theory-based interpretive problems involved in organic chemistry.

Along with the implementation of the system, we have developed the representation of the central knowledge components of the BELIEVER theory. The central common-sense concepts of Person, Plan and Act have been represented as frames. These frames are highly articulated structures which express the core assumptions of the common-sense psychological theory. By expressing these concepts as frames we have been able to provide a representation of these assumptions that can be used to guide and control the overall processes of reasoning about particular persons, plans and actions. The procedural components of the theory have been defined and are closely linked to these frames. This interplay and association between processes and highly articulated structures promises to provide a basis for strongly decomposing the knowledge of the domain. Since the interdependencies of these concepts are represented structurally rather than procedurally, the active database of our MDS-based system provides the basis for communication and cooperation between the processes that monitor these person, plan and act frames.

The definition of these central structural components together with the general system components have also provided a competence theory within which detailed predictions of the BELIEVER theory were specified. These predictions about the structure of summary protocols were tested and borne out by the data. This provides one of the few examples of the verification of predictions derived from work on the development of psychological theory using AI concepts in the process of theory formation.

c) Artificial Intelligence; Representations, Reasoning and Systems Development

Our work in this area continues to be oriented to collaboration with investigators in other Resource projects and to study of basic AI problems that are related to Resource applications. The collaborations involve adaptation and augmentation of existing AI methods and techniques to handle specific key problems identified in the application projects.

The close collaboration with investigators in the Belief Systems area has resulted this year in the development of the AIMDS System for handling problems of action interpretation of the type encountered in the domain of the BELIEVER theory. This system has provided one of the first examples of a working frame-based AI system. In addition, it has led to several important AI results, such as elucidation of the "frame problem" and unification of previous approaches to planning in heuristic problem solving.

Our research in language processing has led this period to two important applications - in Medical Systems and in Belief Systems. In one project, the PEDAGLOT system is being adapted to provide a natural language interface for communicating patient case histories to our glaucoma system. In a second project, PEDAGLOT is providing the basis for implementing the experimental component of a competence theory within which the BELIEVER theory can be evaluated. Empirical work in this area requires the ability to process summaries and other natural language data.

In the basic component of our work on language processing, we continued to develop a language inference system based on a "developmental paradigm" for grammar acquisition. We made progress in the area of coalescing rules of hypothesized grammars, and we started to look into ways of using semantic information to guide the hypothesis formation process.

In another project, which is also focusing on hypothesis formation, we are studying processes of computer assisted acquisition of domain knowledge from empirical data, where knowledge is in the form of weighted production rules. This type of knowledge can be represented as a stochastic graph. This year we obtained several new results in this area. We explored the implications of these results with the help of an experimental program which constructs a stochastic graph from empirical data. Also, we wrote a program which makes use of a file of graph-structured knowledge to make decisions about a domain.

In our work on theory formation in programming, we developed a formation strategy which combines a global, model-guided, approach with a local analysis of special cases. In order to study experimentally this strategy, we are now developing a system for acquiring and handling information about programs in various stages of specification, as well as other knowledge which is relevant to the formation task.

During this period we made important progress in building a strong basis of AI languages for our work. The UCI-LISP and FUZZY programming languages were adapted to the RUTGERS-10 and they were further improved. The availability of these languages made possible the implementation of major parts of AIMDS over a relatively short period of time. Work has now started on exploring the use of

FUZZY (including its features for effective use of incomplete and/or uncertain knowledge) and AIMDS in certain problems of medical decision making.

2. AIM Workshop

The Second AIM Workshop took place June 1 to 4, 1976 near the Rutgers campus, and it was attended by about 150 participants. The program included reviews of recent AI developments in Medicine, Biochemistry and Psychology; lectures and panel discussions on knowledge representation and AI system design; papers summarizing recent AI work in other application areas (outside AIM); and presentations of current research on computer-based biomathematical models. The Workshop included panels on networking and shared resources; in addition, there were a number of informal meetings in which specific projects or issues were discussed in depth. Hands-on experimentation and demonstration of AI systems (which were accessed via TYMNET and ARPANET) were an important feature of the Workshop. All indications are that the Workshop was very effective in stimulating scientific interactions and in disseminating work being done in the area of AIM.

In support of the AIM Workshop series we devoted considerable effort this period to systems development, to related computer and networking enhancements, to preparation of proceedings for the first Workshop, and comprehensive supporting documentation for the second.

A panel on Applications of AI to Science and Medicine was organized for the week following the Second AIM Workshop at the National Computer Conference in New York. It was intended to further augment the dissemination activities of AIM by bringing to a wide audience of professionals in the computer field recent developments in the AIM community.

D) Up-to-Date List of Publications

Amarel, S. and Kulikowski, C. (1972) "Medical Decision Making and Computer Modeling, Proc. of 5th International Conference on Systems Science, Honolulu, January 1972.

Amarel, S. (1974) "Inference of Programs from Sample Computations", Proc. of NATO Advanced Study Institute on Computer Oriented Learning Processes, 1974, Bonas, France.

Amarel, S. (1974) "Computer-Based Modeling and Interpretation in Medicine and Psychology: The Rutgers Research Resource", Proc. on Conference on the Computer as a Research Tool in the Life Sciences", June 1974, Aspen, by FASEB; also appears as Computers in Biomedicine TR-29. June 1974, Rutgers University, also in Computers in Life Sciences. W. Siler and D. Lindberg (eds.), FASEB and Plenum, 1975.

Amarel S. (1976) Abstract of Panel on "AI Applications in Science and Medicine" in 1976 National Computer Conference Program, N.Y., June 7-10, 1976.

Bruce B. (1972) "A Model for Temporal Reference and its Application in a Question Answering Program", in "Artificial Intelligence", Vol. 3, Spring 1972.

- Bruce, B. (1973) "A Logic for Unknown Outcomes", Notre Dame Journal of Formal Logic; also appears as Computers in Biomedicine, TM-35, Nov. 1973, Rutgers University.
- Bruce, B. (1973) "Case Structure Systems", Proc. 3rd International Joint Conference on Artificial Intelligence (IFCAI), August 1973.
- Bruce, B. (1975) "Belief Systems and Language Understanding", Current Trends in the Language Sciences, Sedelow, and Sedelow (eds.) Houton, in press.
- Chokhani, S. and Kulikowski, C.A. (1973) "Process Control Model for the Regulation of Intraocular Pressure and Glaucoma", Proc. IEEE Systems, Man & Cybernetics Conf., Boston, November 1973.
- Chokhani, S. (1975) "On the Interpretation of Biomathematical Models Within a Class of Decision-Making Procedures", Ph.D. Thesis, Rutgers University; also Computers in Biomedicine TR-43, May 1973.
- Fabens, W. (1972) "PEDAGLOT. A Teaching Learning System for Programming Language", Proc. ACM Sigplan Symposium on Pedagogic Languages, January 1972.
- Fabens, W. (1975) "PEDAGLOT and Understanding Natural Language Processing". Proc. of the 13th Annual Meeting of the Asso. of Computational Linguistics, October 30 - Nov. 1, 1975.
- Kulikowski, C.A. and Weiss, S. (1972) "Strategies for Data Base Utilization in Sequential Pattern Recognition", Proc. IEEE Conf. on Decision and Control, Symp. on Adaptive Processes, December 1972.
- Kulikowski C.A. and Weiss, S. (1973) "An Interactive Facility for the Inferential Modeling of Disease", Proc. 7th Annual Princeton Conf. on Information Sciences and Systems, March 1973.
- Kulikowski C.A. (1973) "Theory Formation in Medicine: A Network Structure for Inference", Proc. International Conference on Systems Science, January 1973.
- Kulikowski, C.A. Weiss S. and Safir, A. (1973) "Glaucoma Diagnosis and Therapy by Computer", Proc. Annual Meeting of the Asso. for Research in Vision and Ophthalmology, May 1973.
- Kulikowski, C.A. (1973) "Medical Decision-Making and the Modeling of Disease", Proc. First Internatl. Conf. on Pattern Recognition, October 1973.
- Kulikowski, C.A. (1974) "Computer-Based Medical Consultation - A Representation of Treatment Strategies", Proc. Hawaii Internatl. Conf. on Systems Science, Jan. 1974.
- Kulikowski, C.A. (1974) "A System for Computer-Based Medical Consultation", Proc. Natl. Computer Conf., Chicago, May 1974.
- Kulikowski, C.A. and Safir, A. (1975) "Computer-Based Systems Vision Care", Proceedings IEEE Intercon, April 1975.

- Kulikowski C.A. and Trigoboff, M. (1975) "A Multiple Hypothesis Selection System for Medical Decision-Making", Proc. 8th Hawaii Internatl. Conf. on Systems.
- Kulikowski, C. & N.S. Sridharan, (1975) "Report on the First Annual AIM Workshop on Artificial Intelligence in Medicine. Sigart Newsletter No. 55, December 1975.
- Kulikowski C. (1976) "Computer-Based Consultation Systems as a Teaching Tool in Higher Education, 3rd Annual N.J. Conf. on the use of Computers in Higher Education, March 1976.
- Kulikowski, C., Weiss S., Safir, A. et al (1976) "Glaucoma Diagnosis & Therapy by Computer: A Collaborative Network Approach" Proc. of ARVO, April 1976.
- Kulikowski, C. Weiss, S. Trigoboff, M. Safir, A., (1976) "Clinical Consultation and the Representation of Disease Processes", Some AI Approaches, AISB Conferences, Edinburgh, July 1976.
- LeFaivre, R. and Walker, A. (1975) "Rutgers Research Resource on Computers in Biomedicine, H", Sigart Newsletter No. 54, October 1975.
- LeFaivre, R., (1976) "Procedural Representation in a Fuzzy Problem-Solving System", Proc. Natl. Computer Conf., New York, June 1976.
- LeFaivre, R. (1977) "Fuzzy Representation and Approximate Reasoning", submitted to IJCAI-77, MIT.
- Mathew, R., Kulikowski, C. and Kaplan, K. (1977) "A Multileveled presentation for Knowledge Acquisition in Medical Consultation stems", Proc. MEDINFO 77 (in press).
- Mauriello, D. (1974) "Simulation of Interaction Between Populations in Freshwater Phytoplankton", Ph.D. Thesis, Rutgers University 1974.
- Schmidt, C. (1972) "A comparison of source unidimensional, multidimensional and set theoretic models for the prediction of judgements of trail implication", Proc. Eastern Psych. Asso. Meeting, Boston, April 1972.
- Schmidt, C.F. and D'Addamio, J. (1973) "A Model of the Common Sense Theory of Intension and Personal Causation", Proc. of the 3rd IJCAI, August 1973.
- Schmidt, C.F. and Sedlak, A. (1973) "An Understanding of Social Episodes", Proc. of Symposium on Social Cognition, American Psych. Asso. Convention, Montreal, August 1973.
- Schmidt, C.F. (1975) "Understanding Human Action", Proc. Theoretical Issues in Natural Language Processing: An Interdisciplinary Workshop in Computational Linguistics, Psychology, Artificial Intelligence, Cambridge, Mass., June 1975. Also appears as Computers in Biomedicine, TM-47, June 1975, Rutgers University.

- Schmidt C. (1975) "Understanding Human Action: Recognizing the Motives", Cognition and Social Behavior, J.S. Carroll and J. Payne (eds.), New York: Lawrence Earlbaum Associates, in press. Also appears as Computers in Biomedicine, TR-45, Juhe 1975, Rutgers University.
- Schmidt C.F., Sridharan, N.S., and Goodson, J.L. (1976) Recognizing plans and summarizing actions. Proceedings of the Artificial Intelligence and Simulation of Behavior Conference, University of Edinburgh, Scotland, July 1976.
- Schmidt C. (1976) Understanding human action: Recognizing the plans and motives of other persons. In (eds. J. Carrol and J. Payne) Cognition and Social Behavior, Potomac, Maryland: Lawrence Earlbaum Associates, 1976.
- Schmidt, C.F. and Goodson, J.L. (1976) The Subjective Organization of Summaries of Action Sequences, 17th Annual Meeting of the Psychonomic Society, St. Louis, 1976.
- Sedlak, A.J. (1974) "An Investigation of the Development of the Child's Understanding and Evaluation of the Actions of Others", Ph.D. Thesis, Rutgers University.
- Sridharan, N.S. (1976) "The Frame and Focus Problems in AI: Decision in Relation to the BELIEVER System. Proceedings of the Conference on Artificial Intelligence & the Simulation of Human Behavior, Edinburgh, July 1976.
- Sridharan, N.S. (1976) "An Artificial Intelligence System to Model and Guide Organic Chemical Synthesis, Planning in Chemical Synthesis by Computer, American Chemical Society Press, September 1976.
- Sridharan, N.S. and Schmidt, C.F. (1977), Knowledge-Directed Inference in BELIEVER, Workshop on Pattern-Directed Inference Systems, Hawaii, May 1977.
- Srinivasan, C.V. (1973) "The Architecture of a Coherent Information System: A General Problem Solving System", Proc. of the 3rd IJCAI, August 1973.
- Trigoboff, M. (1976) Propagation of Information in a Semantic Net", Proc. of the Conference on Artificial Intelligence and the Simulation of Behaviour, Edinburgh, Scotland, July 1976; updated version appears in CBM-TM-57, Dept. of Computer Science, Rutgers University, 1977.
- Tucker, S.S. (1974) Cobalt Kinetics in Aquatic Microcosms", Ph.D. thesis, Rutgers University.
- Van der Mude, A. and Walker, A. (1976) "Some Results on the Inference of Stochastic Grammars", abstract in Proc. Symposium on New Directions and Recent Results in Algorithms and Complexity. Dept. of Computer Science, Carnegie-Mellon University.
- Vichnevetsky, R. (1973) "Physical Criteria in the Evaluation of Computer Methods for Partial Differential Equations", Proc. 7th Internatl. AICA Congress, Prague, Sept. 1973; reprinted in Proc. of AICA, Vol. XVI, No. 1, Jan. 1974, European Academic Press, Brussels, Belgium.

- Vichnevetsky, R., Tu, K.W., Steen, J.A. (1974), "Quantitative Error Analysis of Numerical Methods for Partial Differential Equations", Proc. 8th Annual Princeton Conference on Information Science and Systems, Princeton University, March 1974.
- Walker, A. (1975) "Formal Grammars and the Regeneration Capability of Biological Systems", Journal Comp. and Syst. Sciences, Vol. 11, No. 2, 252-261.
- Weiss, S. (1974) "A System for Model-Based Computer-Aided Diagnosis and Therapy", Parts I and II, Ph.D. Thesis, Rutgers University; also Computer in Biomedicine TR-27, Feb. 1974.
- Weiss, S., Kulikowski, C. and Safir, A. (1977) "Glaucoma Consultation Computer", Computers in Biology and Medicine (in press).

II) INTERACTIONS WITH THE SUMEX-AIM RESOURCE

During the past year we have continued to use the SUMEX-AIM resource for program development and testing, for communications between collaborators distributed in different parts of the country and for preparation and running of the AIM Workshop. We continue to access SUMEX-AIM via TYMNET, and to a smaller extent via ARPANET. SUMEX-AIM played a key role in consolidating our network of collaborators in ophthalmology (ONET) and in providing the support needed for establishing a productive collaboration among the ONET investigators. Also, it has been most useful in communicating, planning and helping to set up the information pool for the Second AIM Workshop.

Computing in the Rutgers Research Resource continues to be distributed between SUMEX-AIM and the RUTGERS-10. The two computers are providing complementary resources for our research and for our national collaborations. At present, the distribution of our computing is about 3 to 1 between RUTGERS-10 and

SUMEX-AIM. Our total demand at SUMEX-AIM is estimated at about 5000 connect hours for the current year with most of the work done in INTERLISP (about 80% of our total connect hours) and the rest devoted mainly to communications and to limited program testing within ONET.

The SUMEX-AIM facility was used for demonstrations of AIM programs in first year classes and in second year seminars at the Rutgers Medical School, CMDNJ; CASNET, MYCIN, INTERNIST and PARRY were interactively accessed in these classes and seminars. Another innovative use of SUMEX-AIM has been the collaborative development of the AI HANDBOOK, which is intended to provide a computer-based and network accessible encyclopedic coverage of the AI field for the AIM community and AIM guests. The AI HANDBOOK was initiated by Dr. E. Feigenbaum and his students at Stanford. During the year, a graduate class at Rutgers, given by Dr. S. Amarel, worked on the AI HANDBOOK and contributed several articles.

We find that the SUMEX-AIM bulletin board plays an important role in communicating ideas and information on services among users. Since the MYCIN group at Stanford regularly posts summaries of meetings; and other technical information, on the MYCIN bulletin board, we have been able to keep track of their program and problems. This was particularly useful for our work on IRIS where concepts close to the MYCIN CF formalism are being studied.

6.3 PILOT STANFORD PROJECTS

The following are descriptions of the informal pilot projects currently using the Stanford portion of the SUMEX-AIM resource pending funding, and full review and authorization.

6.3.1 GENETICS APPLICATIONS PROJECT

Computer Science Applications in Genetics

Prof. L. L. Cavalli-Sforza
Department of Genetics
Stanford University School of Medicine

We have been quite satisfied with the use of programs such as REDUCE, MLAB, SPSS. REDUCE has been used by graduate student D. Wagener, to check algebra, and also by L. Cavalli-Sforza and has been of great help in circumstances in which algebraic manipulations were too lengthy for hand verification. Unfortunately REDUCE has a maximum length of algebraic expansions that can be manipulated by computer, which is not always generous enough for our purposes; the maximum allowed was increased but there is now no warning as of when the length of expression overruns the new limits. The penalty is the total loss of the information. If this could be mended, the program would be much more useful. MLAB is very useful for least square fitting of complex systems of equations. SPSS is widely used and well known; it is working fine in the system.

Special modelling efforts involved: 1) a program of information storage and retrieval which may be useful also for analysis of multi-dimensional contingency tables. The material to which it was applied derives from anthropological and archeological survey and excavation data in Calabria, Italy by A. Ammerman. The information collected on coordinates of sites, material found, elevation, land form, soil, ecological and geological data etc. refers to hundreds of sites and will eventually be subject to analysis according to models of growth and spread of Neolithic populations. It is eventually hoped to investigate the power of new techniques of statistical analysis, employing spectral analysis of the matrices representing the data. 2) Similar situations, on the basis of other data available from the literature, are also being investigated by means of simulations of the population growth and spread, e.g. for the Bandkeramik populations in Central Europe. It is thus hoped to obtain, eventually, an explanation of the geographic distribution of genes in Europe, the Middle East and nearby areas, based on the hypothesis that the present distribution reflects predominantly a major radiation of a population of farmers which took place with the spread of agriculture from the Middle East, from 9000 to 5000 years ago. 3) The geographic distribution of genes, as observed today, is analyzed by means of gene frequency maps. We have developed many methods of interpolation of data for map construction, and many methods of graphical display of the maps obtained. We are currently comparing the methods of construction of maps. Some of the methods of construction are fairly sophisticated, but more work will be necessary to develop further our programs so that they can be considered to interpolate intelligently. Our tests of validity are based on eliminating each observation in turn, computing its expected value with the observed one (a sort of jack-knifing). It is clear that results could be improved if this procedure could be carried out simultaneously for several genes and alleles; at the moment it is done for one allele at a time. The simultaneous analysis is an ambitious program but would considerably improve present results. At the moment, for instance, we have no way to make gene frequencies of all alleles at a locus sum to 100% (except approximately, because we cannot consider more than one allele at a

time). In addition, other information on the populations (whether they are isolates, etc.) could be introduced, and verified by the program. Also, specific hypotheses on the evolutionary factors affecting the gene frequencies could be tested more directly. At the moment, the major limitation to these more sophisticated analyses is the availability of computer space.

6.3.2 BAYLOR-METHODIST CEREBROVASCULAR PROJECT

Baylor-Methodist Cerebrovascular Project

John L. Gedye, M.D.
Data Services Research Laboratory
Department of Neurology, Baylor College of Medicine

During the year the Data Services Research Laboratory has had a total of about 2,500 hours of man-effort available, of which about 5% has been devoted to activities directly related to the Sumex pilot study.

I) Summary of research program

A) Technical goals

The general goal of the laboratory - the creation of a computer-based system for the support of clinical research in neurology, as described in the 1975-76 annual report - remains unchanged.

In spite of the limited manpower available during the year, good progress has been made toward the specific goal of developing the PDP11/35-based clinical research system 'CLINSYS' to a point where it can begin to give real support to Departmental projects.

We have made good progress in recent weeks with the development of software which will allow easier access to the resources of SUMEX for users of our local system. It is now possible to give the command 'SUMEX' to our local system executive and have the entire login procedure through to receipt of the "final" SUMEX '@' carried out automatically. Control characters allow the user's terminal to be switched between SUMEX and the local system, and these have been chosen to be compatible with the BANANARD control characters, so that this can be operated without interference.

Facilities have been provided which allow ASCII files to be created on either system and transferred to the other. These facilities will operate under our local PDP11/35 batch system, and we have tested them by creating a test data file of about 1,000 ASCII characters on an account on the PDP11/35, and submitting a batch job (to run at specified time) which logs into SUMEX, transfers the test data file and copies it back again onto the PDP11/35 account and logs out. It then logs in again and repeats the whole process with the latest copy of the file. In this way we hope to estimate the reliability of this form of data transmission - at present it looks as if the error rate will be less than 1 in 16,000 characters - and to lay the foundations for a system that will allow us to make maximum use of SUMEX off-peak time in the projects described below.

B) Medical relevance and collaboration

The development of CLINSYS has continued on the general lines described in the 1975-76 annual report. Specific data acquisition procedures have been designed and implemented for: clinical psychology - both conventional and automated testing techniques have been accommodated; clinical physiology - facilities for the manual entry of Xe133 inhalation regional cerebral blood flow measurements have been provided, and work is now in progress on a system for direct transmission of data to the PDP11/35 from the integral PDP11/05 which is part of the equipment ; and hematology - provision has been made for the acquisition of data from tests of platelet function.

Because of it's central importance, a major emphasis has been placed on making provision for the acquisition of suitably summarised CT scan data, and a number of exploratory studies have been carried out with the result that we hope to have the first edition of a 'CT scan system' working in the near future. This will have an important part to play in future projects.

No further progress has been made with the implementation of a work station incorporating the hand-held OCR wand developed by Recognition Equipment Incorporated - which was described in the 1975-76 report - but we intend to make use of such a 'wand' work station in the context of a system for acquiring data from the radiologist's 'CT scan report' as part of the 'CT' record.

C) Progress summary

The aim of our 'pilot study' remains unchanged - to formulate a project relevant to the activities of the Department which will provide an acceptable and legitimate 'point of entry' for artificial intelligence research, and which will allow the systematic formulation of objectives for the future.

Work has continued along the lines discussed in the 1975-76 report, using, as test data, results from 69 demented patients and 15 controls who had had regional cerebral blood flow measurements. This work has led to a promising 'AI' approach which is now being applied to CT scan data, and when the feasibility of this has been demonstrated the way will be open for work to go head on the implementation of a general purpose program.

D) Publications

There are as yet no publications dealing with the 'pilot study' as such. Certain aspects of the work referred to in this report have been mentioned in publications but these are all currently 'in press'. Details are available on request.

II) Interactions with the SUMEX-AIM resource

- A) Little has so far been achieved by way of collaborations through the network, although the SNDMSG facility has been useful for keeping in touch with contacts made at the 1975 workshop.

It is hoped though, that in the future we may be able to test out the concept of a CT scan archive created by the joint efforts of a dispersed community of users.

- B) For some reason I did not hear about the 1976 workshop until it was over, and so far have heard nothing about a 1977 one. I found the 1975 workshop very useful, and would strongly support the continuation of the workshops in some form - particularly if one could get down to fundamentals with people working on similar problems.

I have kept in close contact with Paul Blackwell at Columbia, Missouri since the 1975 workshop, and we last met at an N.S.F. Conference on 'MATHEMATICAL STRUCTURE IN THE HUMAN SCIENCES' at Penn State in March.

6.3.3 COMPUTER ANALYSIS OF CORONARY ARTERIOGRAMS

Computer Analysis of Coronary Arteriograms

Donald C. Harrison, M.D., Edwin L. Alderman, M.D., and Lynn Quam, Ph.D.
Division of Cardiology, Stanford University Medical School

The goal of this project is to develop computer techniques for automatic acquisition of the anatomic distribution of coronary arteries and a quantitation of the degree of narrowing of these vessels. In order to do this, two different types of image processing techniques will be developed. First, a three-dimensional representation of the coronary arterial tree will be automatically constructed from coronary arteriograms taken sequentially from several different views. Second, the amount of stenosis will be measured by combining information from multiple sequential frames in order to improve resolution and reduce radiographic noise.

BACKGROUND:

Coronary arteriography is the definitive test for the evaluation of patients with coronary artery disease. There is no other test currently available which provides information concerning the location and severity of coronary narrowings and the distribution of coronary blood vessels in the myocardium. Numerous studies document that prognosis in patients with coronary disease reflects the severity of anatomic disease. Coronary vascular anatomy and the extent of lesions are, in an epidemiologic sense, more precise indicators of prognosis than are clinical symptoms.

At the present time, categorization of the extent of coronary vascular disease is based somewhat simplistically on the number of major coronary vessels involved and a rough estimate of the percentage obstruction. Computer representation of the coronary tree, coupled with either interactive or automatic entry of degree of stenosis will permit the development of more precise indices of anatomic disease of the myocardium.

Computer image processing techniques offer the possibility of objectively measuring the severity of coronary stenosis, both at the point of maximal narrowing and averaged over a segment of the vessel.

APPROACH:

An extensive set of image processing functions have been developed and applied to detect the regions of the arteriograms which correspond to the arterial tree. These regions are then transformed to a "skeleton" which roughly corresponds to the midlines of the vessels in the arterial tree. This skeleton is then transformed to a graph representation which can be topologically and geometrically analyzed to distinguish vessel intersections (in the 2-d projection, not real 3-space intersections) from vessel bifurcations. The result is a graph structure interpretation of the arterial tree with quantitation of the

locations (2-d) of bifurcations, and for each vessel segment the path of the vessel midline and the vessel diameter. The computer algorithms are described in more detail in the following sections.

Data Acquisition:

We have digitized a number of 35 mm cine frames from three subjects using both an Optronics film scanner and a Dicomed film digitizer operating at 25 and 50 micron pixel resolution. For each subject frames are manually selected to provide good contrast in the proximal vessels from both LAO and RAO projections and be approximately synchronized within the cardiac cycle.

Pre-processing:

The digitized frames are computer enhanced using high frequency filtering to eliminate the x-ray exposure gradient and emphasize sharp edges which tend to correspond to the vessels.

High contrast areas in the enhanced frames are detected by a simple threshold region detector. Currently, many regions are detected which do not correspond to the arterial tree, but are caused by background features such as vertebra. We are in the process of digitizing another set of frames which have been chosen to include time synchronized pre-injection frames in order to permit background subtraction. The result of this step is a binary image corresponding to high density areas in the frame.

The root of the arterial tree is manually specified by the operator, and a connected point region grower finds all points connected to the root. This usually finds all medium and large sized vessels, and some smaller vessels. Unconnected background is totally eliminated. Sometimes, substantial pieces of the arterial tree are not connected to the root. When this occurs, the operator can run the region grower from new starting points. The result of this step is a binary image corresponding to most of the arterial tree.

We expect that by using background subtraction we can very reliably detect the arterial tree and eliminate most of the manual "hand-holding" in the previous steps.

Arterial Tree Graph Formation:

The binary image of the arterial tree is "skeletonized" by computing the distance transform of the image and connecting peaks and ridges in distance. The distance transform computes for each point in the image, the Euclidean distance to the nearest zero (point not in region). Points at vessel midlines are easily detected because they are local maxima (ridges) in distance from their vessel walls.

The 2-dimensional array of ridge-peak information is next processed to form a graph structure describing the connectivity of vessel segments (distance ridges) to nodes (points where 3 or more ridges converge).

The graph is simplified by detecting and eliminating insignificant terminal segments which are usually the result of noise in the image.

We have now accomplished a significant simplification of the data from the original 2-dimensional array of x-ray density data to an essentially 1-dimensional description of the vessel midlines and points of bifurcation and intersection. This data (when vessel width is included) is sufficient to completely reconstruct the binary image of the arterial tree.

Topologic and Geometric Graph Analysis:

The graph is next analyzed to determine the proximal-distal orientation of each vessel segment. Starting at the distal node of a vessel segment, all segments which are attached to that node must be within 90 degrees in pointing direction. Any segment violating this rule is identified as an intersection. Starting from the root of the arterial tree, all segments are classified by this procedure.

Nodes which have been identified as intersections are now analyzed in order to correspond distal segments with proximal segments according to the a set of rules about arterial topology and geometry.

Having resolved vessel intersections, we now transform the graph to a simple tree structure which corresponds topologically to the arterial tree.

Future Directions:

The above computer algorithms have been successfully applied to the images in a few sets of digitized data. We plan to digitize frames prior to injection to enable background subtraction, which we believe will greatly improve the reliability and accuracy of the initial vessel detection. The algorithms have not yet been tried on cases with abnormal angiograms, and we expect that as more cases are incorporated into our image library, it will be necessary to develop more rules and analytical techniques in order to properly interpret the 2-dimensional images.

Based on the encouraging progress which has been made in processing coronary arteriograms and based on other areas of expertise in image processing within the Stanford University Medical Center, we have developed and submitted on November 1, 1976 to the NHLBI a new grant proposal titled "Computerized Medical Image Processing Laboratory". This proposal contains a detailed report of the progress had been made up to that time and details the further steps which we propose to pursue.

USE OF SUMEX RESOURCE:

Work of this project has been dependent on the SUMEX facility for several reasons. First, this project has not been funded to provide its own computer facilities. Second, although the Stanford Division of Cardiology does have minicomputer systems which could be used for this project, it is considerably

easier to develop image processing and artificial intelligence techniques on a larger scale system in which many powerful tools already exist. It is important in the research phase of this project to be able to easily and quickly perform experiments, without the difficulties of fitting the experimental programs into the small computer memory environment.

6.3.4 QUANTUM CHEMICAL INVESTIGATIONS

Theoretical Investigations of Heme Proteins and Opiate Narcotics

Dr. Gilda Loew
Department of Genetics
Stanford University

SUMEX is used for the calculation of various one-electron electronic properties of iron containing compounds. The programs were formulated and written by David Steinberg, Michael Chadwick and David Lo. David Lo was responsible for converting the program for interactive use on the PDP system. Slight improvements were made by Robert Kirchner and Sheldon Aronowitz has expanded the formulation to include additional spin and oxidation states of the iron atom.

The properties that are calculated include the electric field gradient at the iron nucleus, quadrupole splitting, isotropic and anisotropic hyperfine interaction, spin-orbit coupling and zero field splitting, g values and temperature dependent effective magnetic moments. The calculated values are compared directly to experimental results obtained from published Mossbauer resonance and electron spin resonance spectra. Such a comparison determines not only the reliability with which these properties can be calculated but also gives an indication of the ability of the model of the iron active site to mimic the actual environment found in a particular compound or iron containing protein.

The major input to these properties programs is a description of the electron distribution of the compound under consideration. This description is obtained using a semi-empirical molecular orbital method employing the iterative extended Huckel procedure. Such a calculation requires up to 660K core and is performed elsewhere. When the calculated electron distribution yields a set of calculated properties in agreement with observation, we have increased faith in the description of the model of the active site and can carry the model one step further to make qualitative inferences about certain properties relevant to the biological functioning of the compound.

We are currently performing a systematic study of heme proteins. The electromagnetic properties of these proteins and of synthesized model compounds which mimic the observed behavior of the proteins have been well studied experimentally. Specifically, we have addressed the following problems:

- (1) Cooperativity of oxygen binding to hemoglobin. Calculations have been made for high and low affinity forms of deoxyhemoglobin. This work has been submitted to Nature (Loew and Kirchner).
- (2) The nature of oxygen binding to the heme unit. Calculations were made of model oxyheme compounds with varying oxygen geometry and electron configuration. This work is now in press in the Journal of the American Chemical Society. (Kirchner and Loew).

- (3) The enzymatic cycle of an oxidative metabolizing heme enzyme called cytochrome P-450. This enzyme is responsible for drug metabolism and toxicity and for activation of many chemical carcinogens. Preliminary characterization of the enzymatically active state has been made. This work is in press in the Journal of the American Chemical Society (Loew, Kert Hjelmeland and Kirchner).

In a completely different context, we have been using SUMEX to calculate the conformation of pentapeptides (enkephalins) which have been recently found to be endogenous opiates. The aim of this study is to determine in what way, if any, they can mimic the structure of prototype opiates such as morphine and meperidine. For this work, we use a protein conformation program with empirical interaction potentials. Quantum mechanical conformations calculations of the same peptides are being performed by us elsewhere and the results of the two methods being compared.

6.4 PILOT AIM PROJECTS

The following are descriptions of the informal pilot projects currently using the AIM portion of the SUMEX-AIM resource pending funding, and full review and authorization.

6.4.1 COMMUNICATION ENHANCEMENT PROJECT

Communication Enhancement Project

John B. Eulenberg, Ph.D. and Carl V. Page, Ph.D.
Department of Computer Science
Michigan State University

I) Summary of research program.

A) Technical goals.

The major goal of this research is the design of intelligent speech prostheses for persons who experience severe communication handicaps. Essential subgoals are:

- (1) Design of input devices for persons with greatly restricted movement.
- (2) Development of software for text-to-speech translation.
- (3) Research in knowledge representations for syntax and semantics of spoken English in restricted real world domains.
- (4) Development of micro-computer based portable speech prostheses.

B) Medical Relevance and Collaboration.

We have exchanged visits and had many conversations with Dr. Kenneth Colby of UCLA who is working on similar problems for a domain of people who have aphasia.

The need for such technology in the medical area is very great. Millions of people around the world lead isolated existences unable to communicate because of stroke, traumatic brain injury, cerebral palsy, and other causes. The emergence of inexpensive micro-processors and sound synthesizers makes it possible to develop devices now that can be the prototypes for widespread use.

We have organized institutes to bring together the many professionals who have an interest in this area. Together with the Tufts New England Medical Center, the TRACE Center of the U. of Wisconsin, and the Children's Hospital at Stanford, we have begun the first newsletter for dissemination in this area. Dr. John B. Eulenberg helped to organize the first Federal workshop for governmental agencies who have some interest in funding work in these areas. Represented were the Bureau of Education for the Handicapped, The Veterans Administration, NIMH, NINCDS, NSF, and others. We have also been in touch with United Cerebral Palsy associations at the state and national levels. There is much interest in this area from medical, educational, and governmental communities, but no traditional means of supporting it.

C) Progress summary.

Although some facets of the research have been underway at MSU for several years, we have been using SUMEX-AIM for only six weeks at this time, having received our password in March, 1977. During the last six weeks, we have:

- 1) Designed and built hardware and software allowing us to transmit files to SUMEX from our Nova 2/10 at 300 baud.
- 2) Organized a research team of 4 students possessing background in artificial intelligence led by Dr. Carl V. Page to develop a semantics-based speech generator. We expect to have a prototype running in June (written in SAIL). To this end we are concentrating on semantics associated with personal needs, small talk (weather etc.), and perhaps obtaining geographic directions.
- 3) Have begun conversion of ORTHOPHONE, MSU's large English text-to-speech program from its CDC6500 Fortran implementation to a SAIL version.
- 4) Obtained temporary local support for terminals and tie-lines to use the SUMEX-AIM facility. We requested these in our original proposal but were not granted them. We have to share with others in the use of our tie-lines and terminals. At present the lack of a dedicated tie-line from East Lansing to Tymshare in Ann Arbor or Detroit is a problem for us during 0600 to 0900 PST.

During the past few months, Dr. Richard Reid of our project has:

- 5) Developed a personal communication system for a 10-year-old person who has cerebral palsy. It is micro-computer-based and can accept inputs via an adaptive switch from a series of menus displayed on a TV screen, via Morse code, or by a keyboard. Its outputs can be TV display, hard copy, Morse code, spoken English, Morse code, or musical sounds. We expect to use knowledge gained from the SUMEX-AIM semantics project to specify the content and connection of the choice menus for this project.

During the past three months,

- 6) We have begun to experiment with the interaction of knowledge sources (letter and word frequencies, syntactics, semantics and pragmatics) as a means of anticipating likely inputs and displaying them for a person to choose from.
- 7) Built and tested a myoelectric interface and used it (together with a miniature FM transmitter) for input of changing muscle potentials into a computer. There is reason to believe that this means of input may provide a higher bit rate than any other known means for those people who experience severe motoric problems due to cerebral palsy.

D) Up-to date list of publications. (1976 to date)

For John B. Eulenberg:

"Technical Systems Development, Headend", Interim Report, April, 1976, Experimental Applications of Two-Way Cable Delivery, NSF Grant No. APR 75-14286.

"Interactive New Hired Information Access System with Both Voice and Hard Copy Output: User's Guide to NHQUERRY", April 11, 1976 (With Steven Kludt and Jerome Jackson (Artificial Language Laboratory Report AEB 041176))

"Language Individualization in a Computer-Based Speech Prosthesis System", National Computer Conference, New York, June 9, 1976.

"Individualization in a Speech Prosthesis System", Proceedings of 1976 Conference on Systems and Devices for the Disabled, June 10, 1976.

"The LEAF Language", Interim Report, September, 1976, NSF Grant No. APR 75-14286.

"A Programmable Multi-Channel Modem Output Switch", September 22, 1976, with Joseph C. Gehman and Juha Koljonen (Artificial Language Laboratory Report AEB 092276)

"SMPTE Time Code Interface and Computer-Controlled Video Switcher", with Michael Gorbitt and Dennis Phillips, Interim Report, March, 1977 NSF Grant APR 75-14286.

For Carl V. Page:

"Heuristics for Signature Table Analysis as a Pattern Recognition Technique", IEEE Transactions on Systems, Man and Cybernetics, Vol. SMC-7, No. 2, February 1977.

"Discriminant Grammars, an Alternative to Parsing". with Alan Filipski, Proceedings of the IEEE Workshop on Picture Processing, Computer Graphics, and Pattern Recognition, April 22, 1977.

"Pattern Recognition and Data structures". Chapter in "Data Structures in Computer Graphics and Pattern Recognition" Edited by Allen Klinger, Academic Press, 1977.

During 1976 Dr. Eulenberg presented 15 lectures around the country on his research, was interviewed for TV eight times and was on radio five times.

II) Interactions with the SUMEX-AIM resource.

Again we point out that we have been a part of this community for only about 6 weeks and we will have more to say next year.

A) Examples of medical collaboration and medical use of programs via SUMEX.

The faculty in the MSU College of Human Medicine who teach medical decision making were shown a demonstration of the SUMEX system, MYCIN and PARRY. We plan to present a demonstration to advanced medical students and faculty at the Medical School in the near future.

A member of our Medical School faculty, Dr. Richard Ropple, an expert on myoelectronics, is a member of of our research group.

The Dean of our College of Human Medicine visited our laboratory in April, 1977 and we expect encouragement and collaboration.

B) Examples of sharing, contacts, and cross-fertilization with other SUMEX-AIM projects.

1. We have met with Dr. Kenneth Colby on many occasions including the SUMEX-AIM workshop in June, 1976. Our work in many ways complements his and we have had several worthwhile interchanges of information. We are

converting our major software programs for speech generation and adaptive inputs to the SUMEX AIM system in part so that they can be used by Dr. Colby and his group.

2. Mr. Douglas Appelt, a doctoral student at SU-AI was our principal systems programmer last summer. He is currently doing research in the same area as ours with Dr. Gary Hendrix of SRI. We have used his knowledge of your system (via the message sending routines) to assist us in starting our project. Mr. Appelt will be working with us at MSU again this summer (June-Sept., 1977), and he will be using the SUMEX-AIM system.

6.4.2 AI IN PSYCHOPHARMACOLOGY

Artificial Intelligence in Psychopharmacology

Jon F. Heiser, M.D.
Dept. of Psychiatry and Human Behavior
University of California at Irvine

I. Summary Research Program

A. Technical Goals

1. We propose to construct a computer based system embodying some of the knowledge of an expert in clinical psychopharmacology. Such a system could greatly assist physicians and students who are not specialists in the chemotherapy of mental disorders in choosing the best psychopharmacological treatment for patients for whom such treatment is indicated. The system could also serve as a teaching source of psychodiagnostic and psychopharmacological knowledge.
2. The specific aims of this project are:
 - o To develop a set of MYCIN type rules which are a model of expert clinical teaching, consulting and decision-making for clinical psychopharmacology.
 - o To implement this set of rules in the MYCIN system, and
 - o To evaluate the performance of the resulting system as a teaching and consulting aid.
3. No system currently available or under development approaches the goals of the project in the field of clinical psychopharmacology.
4. It is anticipated that the research will fall into two distinct phases each of approximately 18 months duration. The first and current phase involves evaluating the relevance of the structure of the MYCIN system for use in clinical psychopharmacology by replacing the current infectious disease diagnosis and therapy rules and parameters with psychopharmacology rules and parameters. The second phase will involve accumulating a large body of rules and entering them into the MYCIN system and evaluating their performance. Toward the end of this phase, the behavior of the system will be compared with the behavior of recognized experts working on the Adult Inpatient Psychiatric Service of the UCI Medical Center. This evaluation will focus on the adequacy of the system for representing the knowledge of a skilled psychopharmacologist rather than an actual system performance in the clinical framework.

B. Medical Relevance and Collaboration

1. Medical Relevance

- a. For many years it has been well recognized that potent, effective psychopharmacological agents are frequently used in an unsystematic and irrational manner. The most prescribed medication in the United States today is diazepam (Valium), a minor tranquilizer. The first six most prescribed medications are all psychoactive agents. In California, instances of repetitive use of psychotropic drugs have been reported by 70% of a random sample of adults. About 30% of the sample had used psychotropic drugs in the preceding twelve months. Another study showed that 20% of a medical population was taking psychoactive agents at any given time. These figures do not include alcoholic beverages or non-prescription and illicit drugs with psychoactive properties. Many persons are advised to ingest a daily pharmacologic stew consisting of one or more neuroleptic agents, an antidepressant, an anti-parkinsonian agent, one or more tranquilizers, a hypnotic and possibly a psychostimulant. These regimens are often complicated by non-prescription remedies, alcoholic beverages and illicit drugs. The inevitable drug-drug interactions affect absorption, distribution, binding metabolism and excretion of many drugs.
- b. Each year Americans spend over \$700,000,000 for psychotropic drugs. In a recent year \$150,000,000 was spent on the anti-anxiety agent chlordiazepoxide (Librium). Between 20 and 25 million prescriptions are written each year for diazepam. It is estimated that 170,000,000 prescriptions for psychotropic drugs were written in 1967, and that 202,000,000 prescriptions were written in 1970, more than one for every person in the United States. About 17% of all prescriptions written are for psychoactive drugs. If we include medications in which a psychotropic drug is combined with an antispasmodic vasodilator, or other agent, probably 25% of all prescriptions contain psychotropic drugs. The vast majority of these prescriptions are written by physicians who are not psychiatrists.
- c. Many physicians, including psychiatrists, who are practicing today, completed their formal medical training prior to the 1950's when modern psychopharmacological agents first became available. Their training typically includes no instruction in modern clinical psychopharmacology. Even physicians trained since the mid-1950's cannot be expected to keep abreast of the expanding and changing field of psychopharmacology. The principles and practices recommended a few years ago are rapidly becoming obsolete. A recent study showed that the general knowledge of the pharmacology, physiology, and side effects of psychoactive medications was low in both psychiatrists and non-psychiatrists: less than 20% of the physician subjects were able to devise a psychopharmacologically rational dosage schedule for benzodiazepines. Fifty percent of the non-psychiatrist medical

staff felt that doses up to one gram per day of a tricyclic antidepressant, more than three times the recommended maximum and a potentially fatal amount, might be prescribed for depressive symptoms.

- d. We estimate that there are at least 25 discrete syndromes currently identified in clinical psychiatry, each of which has a unique hierarchy of pharmacological treatment. Each treatment in each section has its own set of potential side effects, adverse reactions and drug-drug, drug-host, drug-age and drug-state of health interaction. In addition, for each therapeutic regimen in each hierarchy, there are several classes of drugs which typically consist of more than one agent or combination of agents which are potentially beneficial and which can be preferentially ranked dependent on several other factors in the clinical situation.

2. Medical Collaboration

1. The principal investigator, Jon F. Heiser, M.D., is a physician who is board certified in psychiatry and in full time teaching, research and University service.
2. Three medical students have participated in this project to date: Clifford Risk, Dana W. Ludwig, and Sue A. Clear.
3. Two resident physicians have participated in this project: Bronco R. Radisavljevic, M.D., and Steven J. Smith, M.D.
4. A Doctor of Pharmacy participates in the research: Pierre J. Menard, Pharm. D.

C. Progress Summary

1. Our initial endeavors to extend a MYCIN-like system to clinical psychopharmacology have been successful. None of the envisioned or predicted problems related to the vagaries of psychiatric terminology and diagnosis have materialized. The system appears to work promptly and accurately to diagnose psychiatric conditions, and in a very rudimentary form, to suggest appropriate pharmacological treatment.

The current system consists of about 20 clinical parameters and about 50 decision rules aimed chiefly at a small number of diagnosed psychiatric conditions.

The following is an example of a sample consultation but realistic case:

Sample Consultation with HEADMED, a Psychopharmacology Advisory System

Patient History

A 37 year old man was brought into the emergency room by his distraught family members. They described how the patient, a well-liked, successful and cheerful but reserved history professor, had become increasingly outspoken, short-tempered, and inattentive to accepted norms of social behavior over the preceding three weeks. An unprovoked argument and physical assault upon a long-time friend motivated the family to seek professional help.

Results of a physical exam and routine laboratory studies were within normal limits. Consultation with the psychiatric service was requested. A first year psychiatric resident re-interviewed the patient and his family, and completed a formal mental status exam prior to obtaining the following consultation from the system.

[Questions are generated by the system and appear in mixed upper and lower case, user responses are in boldface capitals and follow a double asterisk. Added commentary is indented in italics.]

Special options (type for help):

**** AFB**

Hello, Dr. Bowen

Do you wish to run the consultation program

****YES**

29-OCT-76 00:47:26

-----PATIENT-1-----

1) PATIENT'S NAME [The consultation begins with a few standard questions.]

****JOHN DOE**

2) PATIENT'S AGE IN YEARS

**** 37**

3) SEX:

**** MALE**

4) IS John Doe ORIENTED

**** N**

[Since the patient is disoriented, the system begins to Probe a little deeper.]

5) IS John DOE DISORIENTED-TO-PERSON

**** Y**

6) IS John Doe DISORIENTED-TO-PLACE

**** N**

7) IS John Doe DISORIENTED-TO-TIME

**** Y**

8) PLEASE TYPE IN THE ONE WORD WHICH BEST DESCRIBES THE THOUGHT PATTERNS
OR THOUGHT DISORDER OF John Doe

** CONFUSED

CONFUSED not a recognized response

What is THE THOUGHT PATTERN OR THOUGHT DISORDER OF John Doe

Expected responses are: NORMAL PROLIXIC TANGENTIAL CONFABULATION ETC>

Enter HELP for list of user options.

[The initial response was incorrect, so the user was reminded of some of the correct responses. He then requests a display of the entire list. Note that both of these capabilities are part of the core of MYCIN itself. Applying it to this new domain required more than supplying the relevant list of acceptable answers.]

** ?

Expected responses are: NORMAL PROLIXIC TANGENTIAL CONFABULATION
RETARDED LOOSE CONCRETE AUTISTIC DELUSIONAL WORD-SALAD MANIC MUTE
Enter HELP for list of user options.

** TANGENTIAL

9) Has any biological relative of John Doe been evaluated or treated
for a mental or psychiatric problem

** N

10) What is the single word which best describes John Doe's behavior

**

What is the best single word description of John Doe's behavior

Expected responses are: NORMAL AVERAGE UNREMARKABLE CATATONIC-EXCITEMENT
etc.

Enter HELP for list of user options.

** HYPERACTIVE

11) What is the single word which best describes John Doe's mood

** LABILE

12) WHAT IS THE SINGLE WORD WHICH BEST DESCRIBES THE EMOTIONS OR FEELINGS
OF John Doe

** ELATED

13) Does John Doe tend to feel worse or be more symptomatic for the
first several hours after awakening for the day

** WHY

[i.e. WHY is it important to determine whether or not PATIENT-1 tends
to feel worse or be more symptomatic for the first several hours
after awakening for the day]

[4.0] This will aid in determining the DIAGNOSIS OF PATIENT-1.

If

[4.1] John Doe tends to feel worse or be more symptomatic for
the first several hours after awakening for the day

then

there is weakly suggestive evidence (.45) that THE DIAGNOSIS OF
john doe is depression

[RULE017]

[All the standard explanatory capabilities of the core system extend
to the new domain without additional programming effort.]

[back to question 13...]

** N

14) Does John Doe have early morning insomnia

** N

15) Please specify the number of pounds that John Doe lost or gained in the 2 months prior to seeking medical or psychiatric help. Estimate if necessary. Precede weight LOSS with a MINUS sign (for example: -8 - a weight loss of (approximately) 8 pounds within the 2 month period prior to seeking professional help)

** 5

16) Has psychological testing been ordered on John Doe

** N

17) ARE THE PERCEPTIONS OF John Doe NORMAL, i.e., NEITHER DISTORTED NOR DISORDERED

** N

18) DOES John Doe HAVE DEPERSONALIZATION

** N

19) DOES John Doe HAVE EXPERIENCES OF ALIENATION and/or EXPERIENCES OF INFLUENCE

** N

20) DOES John Doe HAVE INCOMPLETE AUDITORY HALLUCINATIONS and/or OTHER HALLUCINATIONS

** Y

21) DOES John Doe HAVE COMPLETE AUDITORY HALLUCINATIONS and/or THOUGHT BROADCASTING

** N

22) what is John Doe's level of consciousness

**

What is THE LEVEL OF CONSCIOUSNESS OF John Doe

Expected responses are: HYPERALERT VIGILANT ALERT NORMAL etc.

Enter HELP for list of user options.

** VIGILANT

23) Is this the first episode of significant mental disturbance that john doe has had

** WHY

[i.e. WHY is it important to determine whether or not this is PATIENT-1's first episode of significant mental disturbance]

[8.0] This will aid in determining THE TREATMENT NAME OF PATIENT-1.

If

[8.1] THE DIAGNOSIS OF John Doe is mania, and

[8.2] this is not John Doe's first episode of significant mental disturbance

then

there is strongly suggestive evidence (.9) that THE TREATMENT NAME OF John Doe is lithium

[RULE071]

[back to question 23...]

[The system has finished with the diagnostic phase, and asks a question to help determine the appropriate therapy.]

** Y

[And finally the system presents its diagnosis and therapy recommendations. The numbers are used to give relative weightings to the conclusions: negative numbers indicate that the evidence tends to rule out that possibility.]

The diagnosis for patient 1 is:

ORGANIC-BRAIN-SYNDROME	(.84)
DEPRESSION	(.37)
MANIA	(.35)
SCHIZOPHRENIA	(.2)
PERSONALITY-DISORDER	(-.4)
NEUROSIS	(-.4)

So the treatment should be

EVALUATION	(.84)
ANTIDEPRESSANT	(.3)
NEUROLEPTIC	(.28)

Recently work has been directed toward expanding the system and revising the representation of psychiatric diagnosis and treatment recommendation.

We have also begun development of a small system to score the Minnesota Multiphasic Personality Inventory (MMPI) psychological test using empirically well established rules easily coded into the MYCIN system.

D. Up to date list of publications.

1. No reports of this work have been published to date.
2. Heiser, J.F. Computer-Aided Diagnosis of Psychiatric Patients. Presented to the Research Meeting, School of Engineering, University of California, Irvine, 7 October 1976.
3. Brooks, R. E. and Heiser, J.F. An Application of Artificial Intelligence to Psychiatry. Presented to:
 - (a) Indian Institute of Technology, Madris, India, 28 September 1976, and
 - (b) Madris Christian College, Madris, India, 3 October 1976.
4. Heiser, J.F. and Brooks, R. E. Artificial Intelligence in Psychopharmacology. Accepted for presentation at the VI World Congress of Psychiatry, Honolulu, Hawaii, 28 August - 3 October 1976.

II. Interactions with the SUMEX-AIM resource

A. Examples of collaboration and medical use of programs via SUMEX

1. As explained fully in the attached research grant application, the MYCIN group has been working informally with Dr. Heiser on the development of a knowledge base of decision criteria for psychopharmacology over the past two years.

B. Examples of sharing, contacts, and cross-fertilization with other SUMEX-AIM projects (via workshops, system facilities, personal contacts, etc.)

1. Dr. Heiser's introduction to the SUMEX-AIM project first occurred at the first AIM workshop held at Rutgers in June 1975.
2. Although Dr. Heiser had previously heard of the MYCIN project, his official collaboration with MYCIN resulted from discussions originating at the first AIM Workshop.
3. A collaborative experiment with Kenneth Mark Colby, M.D., and members of the PARRY project was developed, implemented and analyzed completely on SUMEX-AIM. Enclosed is a rough draft of a paper reporting this "Turing Test" which was performed on-line on SUMEX, with the psychiatrist-judges located at Irvine, the patient-person at UCLA and PARRY at SUMEX.
4. Much technical support has been received freely and continuously from the SUMEX staff and members of the MYCIN team, including basic instruction in the use of SUMEX, TENEX, and MYCIN, principles of knowledge representation in MYCIN, and on-going consultation for details of implementing HEADMED in MYCIN.

Much information has been obtained during three visits to SUMEX and MYCIN, but daily work in this project would be impossible without the ability to converse via links, messages, and telephone conversations with members of the SUMEX and MYCIN staffs.

6.4.3 ORGAN CULTURE PROJECT

Application of Computer Science to Organ Culture

Professor Robert K. Lindsay and Dr. Maija Kibens
The University of Michigan, Ann Arbor

I) Summary of research program

The goal of this research project is to develop new methods for the design and analysis of organ culture experiments, using techniques of artificial intelligence.

The cultivation of organ fragments is an important method for the study of disease processes. In contrast to cell culture, organ culture is designed to inhibit outgrowth of cells and to deal with normal tissue relationships as they exist in the body, divorced from the complexities of organ interaction. The technique involves the maintenance of differentiated cells as a group within their normally associated tissues. With an ability to maintain differentiated tissues in culture, a direct histologic and biochemical assessment of factors influencing an organ is possible. Such a biologic model would permit investigation of the structural and functional effects of various substances directly on the target organ. With a chemically defined medium, the technique would allow a simultaneous evaluation of metabolites or hormones released by the organ fragments.

The research is being done in collaboration with Professors Raymond Kahn, Theodore Fischer, and William Burkel of the Department of Anatomy, the University of Michigan Medical School.

We have been working on methods of image analysis of microscope slides. This has been approached from two directions. On the one hand we are writing programs for special image analysis hardware. These programs will calculate various indices of the condition of the cultivated organ fragments based upon measured morphological features. The second approach is to translate the biologist's verbal descriptions of microscope slides into computer data structures which encode conditions not detectable by our image analysis programs, though readily seen and reported by trained human observers. We have developed a dictionary of anatomical terms and programs for morphological analysis. At present we are working on the syntactic analysis of the scientist's verbal descriptions.

II) Interactions with the SUMEX-AIM resource

We have had valuable contacts with members of the DENDRAL project and the MOLGEN project, which share certain goals and methods with our own work.

6.4.4 NEUROPROSTHESES PROJECT

Neuroprostheses Project

M. G. Mladejovsky, Ph.D., Director
Division of Artificial Organs
University of Utah Medical Center
Salt Lake City, Utah 84112

I. Research Summary

Our research involves the investigation of artificial vision by electrical stimulation of visual cortex and artificial hearing by electrical stimulation of the cochlea. This effort has involved the collaboration of several people from many disciplines, not only from the University of Utah, but also from the Ear Research Institute, Los Angeles; University of Western Ontario, London, Ontario; and Columbia University, New York.

The instrumentation involved is controlled by a minicomputer system consisting of a PDP-8 and a PDP-11/05. Experimental protocols are implemented by programs running in the PDP-11. We sought access to SUMEX in order to use the BLISS-11 compiler which runs on the PDP-10. We are using BLISS-11 as the implementation language for an interactive programming system which will enable more flexible control and variation of our experiments.

The base language we are using is BALM (Malcolm Harrison, "BALM Programmer's Manual", Courant Institute, NYU, 1974). This language is defined in terms of an abstract machine called the MBALM machine. The plan of attack is as follows:

- 1) implement the MBALM machine in BLISS-11
- 2) bring up BALM, using a dummy garbage collector and no virtual memory
- 3) implement garbage collection and virtual memory
- 4) add floating point operations
- 5) add a graphics package
- 6) add real-time capabilities
- 7) provide an interface to PDP-11 machine language

The project has progressed to the point that step 2 is almost complete. This has involved installing a new version of BLISS-11 at SUMEX, writing software to allow file transfers between SUMEX and our PDP-11 (which is connected to the Utah-TIP as a terminal), writing MBALM and various support routines in BLISS-11, implementing an I/O package for BALM in assembly language, and performing a bootstrapping process with the BALM self-definition. Our schedule calls for completing steps 3, 4, and 5 by 1 July 1977. Steps 6 and 7 have not been planned in detail at this time.

We are planning to run the resulting programming system on our PDP-11/05 with 28K core, GT-40 graphics system, and running the RT-11 operating system. Modifying the system to run under a different operating system should be straightforward. However, whether the system will run efficiently on a machine with less than 20K core is questionable. It is too early now to say.

There have been no new publications by our group since our application was filed last year. Currently several papers are in progress but have not yet been submitted for publication. A partial list of previous publications is attached.

When the BALM system has reached a stable state, we will be happy to provide documentation and sources for it to anyone who requests them.

II. Interactions with SUMEX

We have been perfectly satisfied with our use of SUMEX. By far our greatest use of the system has been of text editors and the BLISS-11 compiler.

We have also become acquainted through SUMEX with the OMNIGRAPH graphics package available from NIH and have obtained a copy of the OMNIGRAPH manual. We have not used OMNIGRAPH yet but may wish to in the future. We are considering the features of OMNIGRAPH in the design of the graphics package for our interactive system.

We are quite interested in using the MAINSAIL system being developed at SUMEX and have been told that RT-11 is one of the first operating systems under which it will be available.

IV. Publications

Dobelle, W. H., Mladejovsky, M. G., and Girvin, J.P. Artificial vision for the blind: electrical stimulation of visual cortex offers hope for a functional prosthesis. Science, 183, 1 February 1974, 440-444.

Dobelle, W. H., and Mladejovsky, M. G. Phosphenes produced by electrical stimulation of human occipital cortex and their application to the development of a prosthesis for the blind. J. Physiol., 243, 1974, 553-576.

Dobelle, W. H., Mladejovsky, M. G., Evans, J. R., Roberts, T. S., and Girvin, J. P. 'Braille' reading by a blind volunteer by visual cortex stimulation. Nature, 259, 15 January 1976, 111-112.

Mladejovsky, M. G., Eddington, D. K., Evans, J. R., and Dobelle, W. H. A computer-based brain stimulation system to investigate sensory prostheses for the blind and deaf. IEEE Trans. Biomed. Eng., BME-23, 4 July 1976, 286-296.

Mladejovsky, M. G., Eddington, D. K., Dobelle, W. H., and Brackmann, D. E. Artificial hearing for the deaf by cochlear stimulation: pitch modulation and some parametric thresholds. Transactions of ASAI0, 21, 1975, 1-6.

6.4.5 MATHEMATICAL MODELING OF PHYSIOLOGICAL SYSTEMS

Mathematical Modeling of Physiological Systems

John J. Osborn, M.D., Director
Research Data Facility
The Institutes of Medical Sciences
San Francisco, California 94115

The overall goal of the Institutes of Medical Sciences's collaboration with SUMEX is the application of computer technology to clinical medicine. Our efforts during the past year have been in the fields of knowledge based engineering and mathematical modeling.

We are using our available computer based physiological measurement systems to provide the basis on which physiological interpretation is being developed using knowledge engineering, and to provide the data with which mathematical models are being developed using the SUMEX modeling facility.

BIOMEDICAL KNOWLEDGE ENGINEERING IN CLINICAL MEDICINE (KEMED)

The KEMED system is conceived as an application of the discipline of heuristic based programming to the interpretation of measurements made in clinical medicine. The long range goal of the project is to do research on a biomedical knowledge-based system for interpreting the clinical significance of physiological data. This interpretation will be used to aid in diagnostic decision making and the selection of therapeutic action. Even the best measurements often go unused because of the reasonable reluctance of clinical staff to make measurements whose results they only poorly understand and whose relation to clinical management is ambiguous. We will use techniques of biomedical knowledge engineering to extract and systematize the heuristic knowledge used by experts in the practice of their clinical art. These techniques will be used to construct and utilize a knowledge base to guide inference making by computer programs.

The first program in the KEMED system is designed for interpretation of standard pulmonary function laboratory test data. A knowledge base was developed for interpreting the relationship between measured flows, lung volumes, pulmonary diffusion capacity and pulmonary mechanics and the standard diagnoses of pulmonary function. The knowledge base includes interpretation of measured test results and diagnosis of the type and severity of any pulmonary disease which may be present. The program is being developed as an extension to the MYCIN formalism, and it makes extensive use of the MYCIN structures and programming system. Funding has been requested to continue this work.

MATHEMATICAL MODELING OF PHYSIOLOGICAL SYSTEMS

Mathematical models of the cardio-pulmonary system are being developed to extract clinical physiological information from data acquired by the patient monitoring system. Two approaches are being taken: 1) parsimonious models of the dynamic behavior of CO_2 following an increase in inspired oxygen concentration are being developed for automated patient monitoring application, and 2) a detailed model of the regional behavior of radioactive tracers in the lung is being used as a standard for evaluation of the previous models. The MLAB (Modelling Laboratory) program, available on SUMEX is being used extensively for model development by simulating hypothesized models and for data analysis, i.e., identification of model parameters from experimental data. The CO_2 dilution method has been applied successfully in the ICU and additional funding requested. Two new methods for measuring regional lung function with radioactive tracers have been developed where MLAB was essential and further funding has been requested. MLAB was used to perform an error analysis of the method for measuring regional pulmonary shunt fraction. Also, using MLAB model simulation to understand the complex dynamics of $^{133}\text{-Xenon}$ in the lung-tissue system, a method for measuring intraregional ventilation/perfusion ratio maldistribution has been developed which significantly extends the sensitivity of previous methods. A model of the oculatory system is presently being developed on MLAB in collaboration with the Smith-Kettlewell Institute of the Visual Sciences. We anticipate that their model will be used in the future for treatment of patients with strabismus.

Interface with SUMEX

We use SUMEX through the Tymshare network using a terminal. The text editing facilities of SUMEX, including both text editing and message sending, are excellent additions to our in-house facilities (PDP-11 based system). The message system is particularly useful for communicating ideas and questions with other colleagues using the SUMEX system. Our principal difficulty with SUMEX is turn-around time. Both the MYCIN and MLAB systems are interactive, and the 30-60 second time response times associated with MYCIN and MLAB jobs are at best discouraging.

We have a strong desire to develop in-house capabilities in artificial intelligence. We have already invested significant numbers of hours in developing competence with the MYCIN system, and we are confident of developing an extremely capable staff in heuristic programming. An in-house AI computational capability is a more difficult capability to conceive. Developing

artificial intelligence programming facility on a PDP-11 based system remains a significant long-term interest. The satellite capability offers both the potential of not continuing to provide additional load on SUMEX, and it offers the potential of more rapid interaction with the user.

The SUMEX facility contributed to the following grant applications and articles:

Bibliography

- 1) Simulation to Relate Measured Gas Concentrations at the Mouth to Pulmonary Mechanics and Perfusion. J.C. Kunz, R.R. Mitchell, D.H. McClung, J.J. Osborn, Submitted to the 1977 ACEMB.
- 2) Identifiability of Pulmonary and Recirculation Parameters Following Sequential Bolus Inputs of ^{133}Xe . R.R. Mitchell, R.J. Fallat. Submitted to the 1977 ACEMB.
- 3) Simulation of Intraregional Ventilation-Perfusion Ratio Mal-distribution. J.C. Glaub, R.R. Mitchell, R.J. Fallat. Submitted to the 1977 ACEMB.
- 4) Measurement of Residual Volume and Ventilation Distribution Using Helium and a Five Vital Capacity Breath Maneuver. R.R. Mitchell, Technical Report 32, Institutes of Medical Sciences, Feb. 1977.
- 5) Identification of Human Oculomotor System Parameters with Application to Strabismus. N.K. Gupta, A.V. Phatak, Systems Control; R.R. Mitchell, Heart Research Institute and Carter Collins, Smith-Kettlewell Institute, Institutes of Medical Sciences. Submitted to Joint Automatic Control Conference, 197..

6.4.6 PUFF/VM PROJECT

PUFF/VM - Pulmonary Function and Ventilator Management Project

John J. Osborn, M.D.
The Institutes of Medical Sciences (San Francisco)

and

E. A. Feigenbaum
Computer Science Department, Stanford University

Note: The PUFF/VM project is the outgrowth of the efforts of Prof. Feigenbaum's group at Stanford to establish new applications areas for AI in medical research. It represents a collaboration with Dr. Osborn's group which has been working on another AIM pilot project titled "Mathematical Modeling of Physiological Systems". A PUFF/VM proposal is currently pending with NIH and PUFF/VM is being reviewed in parallel by the AIM Executive Committee for separate pilot status.

1. Overall Objectives:

Our immediate objective is to develop a computer programming system for interpreting the clinical significance of measures of pulmonary function. We hope to develop this system for diagnostic use in the pulmonary function laboratory and to aid diagnosis and ventilator management of respiratory insufficiency in the intensive care unit. We hope to demonstrate the clinical effectiveness of such a system for improving the accuracy and timeliness of diagnosis.

Our long range goal is to develop an integrated system for making and interpreting measures of pulmonary function. We believe that this is possible because of the present and potential contribution of instrumentation and data analysis systems to the diagnosis and clinical management of pulmonary distress. We believe, in addition, that the discipline of knowledge-based heuristic programming is potentially the best basis on which to develop a system for automatically interpreting the results of the measures of pulmonary function.

We aim, in the long run, to develop an inexpensive enough implementation that the system will find wide acceptability in the delivery of clinical care.

[Further details will be furnished by Dr. Feigenbaum on request]

Appendix I

OVERVIEW OF ARTIFICIAL INTELLIGENCE RESEARCH

ARTIFICIAL INTELLIGENCE RESEARCH

What is it? What has it achieved? Where is it going?

Excerpt from a report by
Professor Edward A. Feigenbaum
Stanford University

May 1975

[The contents of this appendix have been deleted for brevity -
copies may be had upon request to Dr. Lederberg]

Appendix IIAI HANDBOOK OUTLINENOTE:

The following material is a tentative outline of a handbook on artificial intelligence planned for publication. It is not to be cited or quoted out of the context of this report without the express permission of Professor E. A. Feigenbaum of Stanford University. This handbook is intended for two kinds of audience; computer science students interested in learning more about artificial intelligence, and engineers in search of techniques and ideas that might prove useful in applications programs. Articles in the first seven sections are expected to appear in the first volume to be published in preliminary form by September 1977. The remaining articles are expected to appear in the second volume to be published in preliminary form by June 1978.

The following is a brief checklist that was used to guide the computer science students engaged in writing articles for the handbook. It is, of course, only a suggested list.

- i) Start with 1-2 paragraphs on the central idea or concept of the article. Answer the question "what is the key idea?"
- ii) Give a brief history of the invention of the idea, and its use in A.I.
- iii) Give a more detailed technical description of the idea, its implementations in the past, and the results of any experiments with it. Try to answer the question "How to do it?".
- iv) Make tentative conclusions about the utility and limitations of the idea if appropriate.
- v) Give a list of suitable references.
- vi) Give a small set of pointers to related concepts (general/overview articles, specific applications, etc.)
- vii) When referring in the text of an article to a term which is the subject of another handbook article, surround the term by +'s; e.g. +Production Systems+.

AI Handbook Articles

I. INTRODUCTION

- A. Philosophy
- B. Relationship to Society
- C. History
- D. Conferences and Publications

II. HEURISTIC SEARCH

- A. Heuristic Search Overview
- B. Search Spaces
 - 0. Overview
 - 1. State-space representation
 - 2. state-space search
 - 3. Problem-reduction representation
 - 4. AND-OR trees and graphs
- C. "Blind" Search Strategies
 - 1. Overview
 - 2. Breadth-first searching
 - 3. Depth-first searching
 - 4. Bi-directional searching
 - 5. Minimaxing
 - 6. Alpha-Beta searching
- D. Using Heuristics to Improve the Search
 - 1. Overview
 - 2. Best-first searching
 - 3. Hill climbing
 - 4. Means-ends analysis
 - 5. Hierarchical search, planning in abstract spaces
 - 6. Branch and bound searching
 - 7. Band-width searching
- E. Programs employing (based on) heuristic search
 - 1. Overview
 - 2. Historically important problem solvers
 - a) GPS
 - b) Strips
 - c) Gelernter's Geom. Program

III. AI Languages

- A. Early list-processing languages
- B. Language/system features
 - 0. Overview of current LP languages
 - 1. Control structures
 - 2. Data Structures (lists, associations,
 - 3. Pattern Matching in AI languages
 - 4. Deductive mechanisms
- C. Current languages/systems
 - 1. LISP, the basic idea
 - 2. INTERLISP
 - 3. QLISP (mention QA4)
 - 4. SAIL/LEAP
 - 5. PLANNER
 - 6. CONNIVER
 - 7. SLIP
 - 8. POP-2
 - 9. SNOBOL
 - 10. QA3/PROLOGUE

IV. Representation of Knowledge

A. Overviews

1. Survey of representation techniques
2. Issues and problems in representation theory

B. Representation Schemes

1. Predicate calculus
2. Semantic nets -- Quillian, Hendrix, LNR
3. Production rules
4. MERLIN
5. Procedures (SHRDLU, actors, demons)
6. Frames
7. Componential analysis
8. Scripts
9. KRL
10. Multiple Knowledge sources - Blackboard
11. Query languages
12. FOL

V. SPEECH UNDERSTANDING SYSTEMS

A. Overview (include a mention of ac. proc.)

B. Integration of Multiple Sources of Knowledge

C. The ARPA speech systems

1. HEARSAY I
2. HEARSAY II
3. SPEECHLIS
4. SDC-SRI System (VDM3)
5. DRAGON

VI. Natural Language

A. Overview - History & Issues

B. Representation of Meaning

C. Grammars and Parsing

1. Review of formal grammars
2. Extended grammars
 - a. Transformational grammars
 - b. Systemic grammars
 - c. Case Grammars
3. Parsing techniques
 - a. Overview of parsing techniques
 - b. Augmented transition nets, Woods
 - c. CHARTS - GSP

D. Text Generating systems

E. Machine Translation

1. Overview & history
2. Wilks' machine translation work

F. Famous Natural Language systems

1. Early NL systems (SAD-SAM through ELIZA)
2. PARRY
3. MARGIE
4. LUNAR
5. SHRDLU, Winograd

VII. Applications-oriented AI research (overview)

- A. Chemistry
 - 1. Mass spectrometry - DENDRAL
 - 2. Organic Synthesis - overview
- B. Medicine
 - 1. MYCIN
 - 2. Others
- C. Psychology and Psychiatry
 - 1. Protocol Analysis (Waterman and Newell)
- D. Math systems
 - 1. REDUCE
 - 2. MACSYMA (mention SAINT)
- E. Business and Management Science Applications
 - 1. Assembly line/ power distrib.
- F. Miscellaneous
 - 1. LUNAR
 - 2. Education
 - 3. SCHOLAR
 - 4. SOPHIE
 - 5. SRI computer-based consultation
 - 6. RAND--RITA production rule system
 - 7. Randevous - Query languages

VIII. AUTOMATIC PROGRAMMING

- A. Overview
- B. Program Specification Techniques
- C. Program Synthesis techniques
 - 0. Overview
 - 1. Traces
 - 2. Examples
 - 3. Problem solving applications to AP
 - a. Sussman's Hacker
 - b. Program Synthesis by Theorem Proving
 - 4. Codification of Programming Knowledge
 - 5. Integrated AP Systems
- D. Program optimization techniques
- E. Programmer's aids
- F. Program verification

IX. THEOREM PROVING

- A. Overview
- B. Resolution Theorem Proving
 - 1. Basic resolution method
 - 2. Syntactic ordering strategies
 - 3. Semantic & syntactic refinement
- C. Non-resolution theorem proving
 - 0. Overview
 - 1. Natural deduction
 - 2. Boyer-Moore
 - 3. LCF
- D. Uses of theorem proving
 - 1. Use in question answering

- 2. Use in problem solving
- 3. Theorem Proving languages
- 4. Man-machine theorem proving
- E. Predicate Calculus
- F. Proof checkers

X. Human Information Processing -- Psychology

- A. Perception
- B. Memory and Learning
 - 1. Basic structures and processes in IPP
 - 2. Memory Models
 - a. semantic net memory models
 - b. HAM (Anderson & Bower)
 - c. EPAM
 - d. Productions (HPS)
 - e. Conceptual Dependency
- C. Psycholinguistics
- D. Human Problem Solving
 - 0. Overview
 - 1. PBG's
 - 2. Human chess problem solving
- E. Behavioral Modeling
 - 1. Belief Systems
 - 2. Conversational Postulates (Grice, TW)
 - 3. PARRY

XI. VISION

- A. Overview
- B. Polyhedral or Blocks World Vision
 - 1. Overview
 - 2. Guzman
 - 3. Falk
 - 4. Waltz
- C. Scene Analysis
 - 1. Overview
 - 2. Template Matching
 - 3. Edge Detection
 - 4. Homogeneous Coordinates
 - 5. Line Description
 - 6. Noise Removal
 - 7. Shape Description
 - 8. Region Growing (Yakamovsky, Olander)
 - 9. Contour Following
 - 10. Spatial Filtering
 - 11. Front End Particulars
 - 12. Syntactic Methods
 - 13. Descriptive Methods
- D. Robot and Industrial Vision Systems
 - 1. Overview and State of the Art
 - 2. Hardware
- E. Pattern Recognition
 - 1. Overview
 - 2. Statistical Methods and Applications

- 3. Descriptive Methods and Applications
- F. Miscellaneous
 - 1. Multisensory Images
 - 2. Perceptrons

XII. ROBOTICS

- A. Overview
- B. Robot Planning and Problem Solving
- C. Arms
- D. Present Day Industrial Robots
- E. Robotics Programming Languages

XIII. Learning and Inductive Inference

- A. Overview
- B. Samuel Checker program
- C. Winston -- concept formation
- D. Pattern extrapolation problems--Simon,
- E. Overview of Induction
- F. AQVAL (Michalski at U.Ill)
- G. Parameter adjustment of linear functions
- H. Rote learning
- I. D.A. Waterman's machine learning of heuristics
- J. Learning by debugging
- K. Learning by parameter Adaptation
- L. Signature & move phase tables

XIV. Reasoning and Planning

- A. Reasoning by analogy
 - 1. Overview
 - 2. ZORBA
- B. planning
 - 1. NOAH
 - 2. ABSTRIPS

Appendix IIISUMMARY OF MAINSAIL LANGUAGE FEATURES

MAINSAIL LANGUAGE FEATURES

Clark R. Wilcox
Stanford University

Portable ALGOL-like language with dynamic memory support

MAINSAIL is an ALGOL-like language with dynamic memory support for strings, arrays, records, modules and files. The driving force behind its design is that it provide for the development of portable software. At the same time, low-level features allow the programmer to deal with the underlying representation of data aggregates. These low-level features have made it possible for most of the runtime system to be written as MAINSAIL modules.

Intended applications

MAINSAIL is not oriented toward any particular application. The flexible use of memory makes it suitable for tasks with memory requirements which are difficult to predict prior to execution, as is often the case with knowledge representation. The string capabilities facilitate word processing applications such as compilers, text editors and document preparation, and "friendly" interactive programming. These same facilities require runtime support, so that a MAINSAIL program is not a stand-alone body of code, and thus may not be appropriate for some primitive system utilities.

Portability

A primary goal is that compatible implementations be provided on a variety of computer systems. Programs which are written for portability should be able to execute on any of the implementations with the same effect. Such programs must adhere to reasonable constraints with regard to data and memory ranges, as described in the language manual. Programs which violate these constraints are not considered portable, and thus may behave differently on different implementations. This design for portability raises a number of questions with regard to how well MAINSAIL will fit any particular machine. It is too early to provide a conclusive answer to such concerns, though it appears that many machines will efficiently support MAINSAIL implementations.

Modularity

In addition to the more obvious effects the machine-independent design has on data types and operations, it also necessitates a model of runtime interactions which can be supported on a broad range of computers. In particular MAINSAIL must be able to execute in a limited address space, which means that programs must be broken into pieces (modules) which need be in memory only when executing. The inability to characterize linkage and overlay systems in a

machine-independent manner has forced MAINSAIL to take over these functions, and thus assume duties often considered part of the operating system.

A MAINSAIL program consists of an open-ended collection of modules, i.e., the programmer need not specify what modules make up a program. The modules may originate from many files at execution, as contrasted to the common approach of having a single "save file" or "load module" which may contain an overlay structure.

The modules are compiled separately and assembled into a form which does not require linkage prior to execution. MAINSAIL resolves all inter-module references at runtime. Modules are automatically brought into memory as needed. If there is insufficient room in memory for an incoming module, MAINSAIL automatically swaps out one or more resident modules to make room. This swapping could involve i/o to an external device or memory mapping. Modules are position-independent, i.e., they do not contain references to fixed memory locations. Thus they may be moved about during execution, and need not be swapped into the same memory locations from which they were swapped out. This generalization of the traditional overlay structure will make possible the implementation of sizeable programs in a limited address space, while at the same time utilizing the minimum possible memory on larger systems.

Range of data types

In order to allow efficient operation on machines with a small word size, yet access to large values when necessary, MAINSAIL offers both short and "long" data types: integer, long integer, real, long real, bits and long bits. In practice the long forms are used much less frequently than the short forms, and thus can be simulated if necessary with no major degradation in efficiency. These data ranges have been chosen to fit the range of machines for which MAINSAIL is intended.

Strings

A MAINSAIL string is a variable length sequence of characters. The programmer does not need to specify a maximum length for a string as is common in many languages. Instead, MAINSAIL keeps track of the current number of characters in a string and automatically handles storage allocation. Most existing general-purpose languages have omitted a full implementation of strings, apparently under the assumption that they could not be efficiently implemented, and were dispensable. However, the hardware design trend is toward microprogrammed instruction sets which support string operations, in view of the increasing acceptance of computers for word-processing.

Classes, records and pointers

MAINSAIL employs a general notion of "class" as a collection of data and procedures fields. Classes serve two purposes: they specify the interfaces through which modules communicate with one another; and they are used as templates for the creation of and access to records. A record is a dynamically allocated memory area which contains data corresponding to the fields of the class to which it belongs. The fields of a record are accessed by means of a pointer to the record, combined with the name of the field. The pointer must have been associated with the record's class when it was declared.

The notion of "prefix class" was introduced to provide for a hierarchy of classes. A class which is declared with a prefix class is conceptually made a member of the prefix class, and inherits the fields of the prefix class as its initial fields. For example, the concept "doubly-linked list" may be represented as a class with two pointer fields, say "left" and "right". Any other class will automatically inherit these two fields if it is defined as a doubly-linked-list class.

The language contains rules which govern the use of pointers according to the relationships between classes and prefix classes. MAINSAIL provides for secure use of pointers in the majority of cases, but allows insecure operations if desired.

Arrays

MAINSAIL's implementation of arrays is quite flexible in that it allows the programmer full control over the creation and disposal of arrays. This is to be contrasted with classical ALGOL, where array allocation is tied to block structure. An array is actually a pointer to a record, and thus is allowed many of the same constructs provided for pointers, such as assignment, equality comparison, and parameter passing. An array may be a field of a class, so that any number of records may be allocated which contain array fields. This capability is particularly useful in image processing, where flexible array allocation can significantly simplify program logic.

Procedures

Procedures play a major role in MAINSAIL. Procedures may be typed for use in expressions. There are three simple parameter passing mechanisms: USES passes the value; PRODUCES passes a value back to the caller; and MODIFIES passes and returns a value. Optional arguments, repeatable arguments, and generic procedures provide useful syntactic constructs. Any procedure may be invoked recursively. Other procedure characteristics are COMPILETIME (if all arguments are constants, the procedure is evaluated during compilation), INLINE (produces "in-line" code), and CODED (supports assembly language coding).

Embedded assembly language

A number of facilities support the use of assembly language within a MAINSAIL program: CODED procedures, the Code statement, and the various forms of encoding variable offsets. Of course assembly language cannot appear within a machine-independent program, but nevertheless there are many instances when the target machine is known. The MAINSAIL interface to each operating system makes extensive use of the assembly language facilities.

Compiletime support

Most present-day compilers were designed to work in a sequential access mode, and suffer from the resulting limitations. The MAINSAIL compiler was designed with the understanding that the source files would be on random-access devices, so that it need not progress through the file in a strictly linear fashion. Any number of nested input files are allowed, in fact the same file may be scanned several times during compilation (contrast this with a compiler designed for input from punched card decks).

Compilation involves interaction with the user in that the programmer can put messages in a source file which are displayed during compilation. The user can affect the course of the compilation by specifying the names of files to be compiled as requested by directives within the file being compiled, and by defining values which govern the scanning of the source text. The compiler has the ability to quickly search through a file for the text to be compiled as specified either by earlier source text, or interactively by the user. This allows a single file to be made a repository of fragments of source text needed during many different compilations, and quickly searched during a particular compilation.

Conditional compilation allows an arbitrarily complicated expression (ultimately made up of constant operands) to be evaluated by the compiler to determine whether a particular segment of the source file is to be ignored. In general, the compiler will evaluate all expressions involving only constant operands (of type boolean, (long) integer, (long) bits, and string) and compiletime procedures. These facilities are quite important when building a large parameterized system.

A save and restore facility allows the current state of the symbol table to be saved. It may be restored during a later compilation to avoid recompiling unchanged text. This is particularly useful for the development of a collection of modules all of which utilize one or more common "header" files.

A comprehensive macro facility provides for the definition of constants, arbitrary text, and arbitrary text with parameters. Many commonly used constants are predefined, especially as needed by the system procedures to simplify passing of bits parameters consisting of predefined "flags".

File system

A simple yet powerful file system has been designed which, like all features of MAINSAIL, is guaranteed on every implementation. When a file is opened for use, the program specifies whether it contains text or data (binary), and whether access is sequential or random. A fundamental assumption is the ability to communicate with a controlling terminal, called the tty ("teletype"). For example, error messages are output to tty, and a response is expected.

Appendix IVMICROPROGRAMMED MAINSAIL PLANS

Plans for a Microprogrammed Implementation of MAINSAIL

Clark R. Wilcox
Stanford University

In this appendix we shall discuss our plans for a microprogrammed implementation of MAINSAIL. The goal of this research is to determine the feasibility of distributing a cost-effective integrated hardware-software programming environment. A computer which operates under the control of a microprogrammable control store offers a new approach to efficient program execution which we summarize below. We feel this approach could offer the means of developing reasonably-priced computing resources with the capability of executing programs which are too demanding for present mini-computers. It appears that such machines may be widely available within a few years. We propose to purchase the necessary hardware to enable us to develop a microprogrammed MAINSAIL implementation.

The emulation approach to high-level language implementation

Traditional implementations of high-level language involve translation to the fixed machine languages of the target machines. Such machine languages have not been designed for the efficient representation of high-level languages, with the result that an excessive number of overhead instructions are required to map the high-level language into its directly-executable machine code "surrogate". With the advent of microprogrammable computers with writable control stores, a different approach appears to have great promise for the efficient execution of high-level languages.

A micro-coded computer executes the instructions in main memory under control of the micro program. Thus the machine code may be viewed as data which is interpreted, or emulated, by the micro program, rather than as direct signals to the hardware. The micro program is written in a more primitive machine code called micro code, which (usually) directly controls the hardware. Most micro-coded computers have been designed for the emulation of a particular machine code, and thus the micro-code is simply a means of reducing the complexity of the hardware while perhaps providing a "higher-level" machine code. The micro-code is placed into a high-speed memory (relative to main memory), so that many micro instructions can be executed in the time it takes to fetch a single instruction from main memory.

The same technique of interpreting a particular machine code with a micro program can be broadened to the ability to interpret an arbitrary machine code. Such a micro computer is called a "soft" machine, or "universal host", since it is not oriented toward any particular machine code. Instead, the language implementor chooses a suitable machine-code representation. A compiler is constructed which translates into this representation, and a micro program is

written which interprets the representation. This approach is known as a "directly executable language", or DEL, since the high-level language has been translated into a form tailor-made for it. The unnecessary overhead instructions are eliminated, with a resulting decrease in program representation and increase in execution speed. There is evidence [3,4,6] that this approach can provide substantial dividends.

A MAINSAIL Directly Executable Language (DEL)

We propose to design a MAINSAIL DEL and implement it on a microprogrammable computer. The goal is to evaluate the economic and technical advantages of exporting a combined hardware-software environment for program development and distribution. In particular, we want to orient MAINSAIL's design and implementation toward such an emulation approach and compare the resulting "MAINSAIL machine" with conventional implementations.

We are interested in determining whether a "soft" machine of this sort can be provided cheaply enough to serve as a basis for the distribution of software which presently requires expensive hardware facilities. Hardware which can be specifically tailored for high-level language execution may provide the quickest route to the economically viable distribution of programs which exceed the limits of present general-purpose mini-computers.

This work will complement the on-going implementations of MAINSAIL on conventional hardware. Thus we will be in a unique position to compare the two approaches. We expect the MAINSAIL DEL to outperform other MAINSAIL implementations in much the same way that DELtran (a DEL for FORTRAN II) outperforms FORTRAN II [3]. Initial measurements show that the DELtran representation is less than one fifth the size of the code generated by the FORTRAN-H optimizing compiler, and executes about five times faster.

MAINSAIL is perhaps better suited to the emulation approach than FORTRAN because of the locality of reference provided by procedures, records and modules. A preliminary DEL has already been designed for MAINSAIL, but further work is necessary before we can predict (or demonstrate) size and execution comparisons with standard implementations. There is much work to be done in determining the efficient representation of ALGOL-like languages for the purpose of emulation, and providing data from actual implementations.

A MAINSAIL DEL could provide facilities which are impossible to provide in an efficient manner on conventional machines. These facilities relate to the monitoring of the program during execution. Since the emulator is simply a program written in micro code, it can be made to perform any kind of execution-time checks with no need to alter the DEL. By contrast, the MAINSAIL compiler must generate different code depending on the amount of checking to be performed.

The emulator can also provide execution profiles and comprehensive debugging facilities such as instruction traps and single stepping. We expect to provide several emulators which are oriented toward particular types of execution, e.g. a "fast" emulator which maximizes execution speed, a "careful" emulator which provides comprehensive runtime checks, a "performance monitoring" emulator which gathers information concerning program execution, and a "debugging" emulator which allows interactive debugging.

Another advantage of the emulation approach is the simplifications in the compiler. Since the compiler will translate MAINSAIL to its own DEL, the code generators become almost trivial. MAINSAIL operations which require many instructions on existing machines can be compactly represented with the DEL. The compiler need not worry about register optimization since there will be no registers in the DEL representation. Since the MAINSAIL DEL is a close representation of the source code, there is no reason to "drop into assembly language" since any "sensible" program which could be written in the DEL could more easily be written in MAINSAIL.

Hardware support

To support this development, we propose the purchase of a dynamically micro-programmable machine with such supporting hardware as is necessary. This machine should be a universal host in the sense that it is not already oriented towards a particular machine code. Its software support is of little consequence since we will design our own operating system and high-level language support.

We are interested in implementing sophisticated programs, and thus require a large address space (say 24 bits) and 32-bit arithmetic. We need sufficient control store, say 16K words, to support a debugging emulator and selected parts of the operating system. The micro store must be able to quickly transfer words to and from main memory, in particular we want to be able to quickly switch emulators. There must be facilities for interface to a variety of peripherals, and to other computers.

There are some machines now available along these general lines (e.g. [1]), with the introduction of more imminent. Indeed, manufacturers are beginning to include user-microprogrammable features with new models of their traditional hardware, e.g. Digital Equipment Corporation's PDP-11/60 and Data General's Eclipse.

One such machine, EMMY, has been developed by the Stanford Emulation Laboratory, under the direction of Professor Michael Flynn of the Department of Electrical Engineering [2,5]. EMMY is a universal host machine which closely fits our needs. It is an unbiased yet efficient host for a wide range of target machine architectures. EMMY is scheduled to go into production in late 1977 by ICL of England (the emulation laboratory has been involved in the development of a prototype). We feel that this machine would suit our needs, but further evaluation is necessary.

We expect most of development of the MAINSAIL DEL to be independent of any particular micro program representation. In particular, we are not at this time proposing to carry out any hardware design to orient the host processor towards MAINSAIL, though this approach would be reasonable if a large number of processors were to be distributed solely to support MAINSAIL execution.

References

1. Burroughs Corp., "B-1700 Systems Reference Manual," Burroughs Corp., Detroit, Michigan, 1972.
2. Flynn, M. J., Hoevel, L. W., and Neuhauser, C. J., "The Stanford Emulation Laboratory," Digital Systems Lab., Technical Report No. 118, Stanford University, June 1976.
3. Hoevel, L. W. and Flynn, M. J., "The Structure of Directly Executed Languages: A New Theory of Interpretive System Support," Digital Systems Laboratory, Technical Report No. 130, Stanford University, March 1977.
4. Hoevel, L. W., "DELtran Principles of Operation," Digital Systems Laboratory, Technical Note No. 108, Stanford University, March 1977.
5. Neuhauser, C. J., "An Emulation Oriented, Dynamic Microprogrammable Processor," Digital Systems Lab., Technical Note No. 65, Stanford University, October 1975.
6. Wilner, W., "Burroughs B-1700 Memory Utilization," AFIPS Proceedings, Vol. 41-I, FJCC, 1972, pp. 579-586.

Appendix VAIM MANAGEMENT COMMITTEE MEMBERSHIP

The following are the membership lists of the various SUMEX-AIM management committees at the present time:

AIM EXECUTIVE COMMITTEE:

LEDERBERG, Joshua, Ph.D. (Chairman)
Department of Genetics, S331
Stanford University Medical Center
Stanford, California 94305
(415) 497-5801

AMAREL, Saul, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-3546

BAKER, William R., Jr., Ph.D. (Executive Secretary)
Biotechnology Resources Program
National Institutes of Health
Building 31, Room 5B43
9000 Rockville Pike
Bethesda, Maryland 20014
(301) 496-5411

LINDBERG, Donald, M.D. (Adv Grp Member)
605 Lewis Hall
University of Missouri
Columbia, Missouri 65201
(314) 882-6966

MYERS, Jack D., M.D.
School of Medicine
Scaife Hall, 1291
University of Pittsburgh
Pittsburgh, Pennsylvania 15261
(412) 624-2649

AIM ADVISORY GROUP:

LINDBERG, Donald, M.D. (Chairman)
605 Lewis Hall
University of Missouri
Columbia, Missouri 65201
(314) 882-6966

AMAREL, Saul, Ph.D.
Department of Computer Science
Rutgers University
New Brunswick, New Jersey 08903
(201) 932-3546

BAKER, William R., Jr., Ph.D. (Executive Secretary)
Biotechnology Resources Program
National Institutes of Health
Building 31, Room 5B43
9000 Rockville Pike
Bethesda, Maryland 20014
(301) 496-5411

BOBROW, Daniel G., Ph.D. [Term expiring]
Xerox Palo Alto Research Center
3333 Coyote Hill Road
Palo Alto, California 94304
(415) 494-4438

FEIGENBAUM, Edward, Ph.D.
Department of Computer Science
Polya Hall, Room 213
Stanford University
Stanford, California 94305
(415) 497-4079

FELDMAN, Jerome, Ph.D. [Term expiring]
Department of Computer Science
University of Rochester
Rochester, New York
(716) 275-5671

LEDERBERG, Joshua, Ph.D. (Ex-officio)
Principal Investigator - SUMEX
Department of Genetics, S331
Stanford University Medical Center
Stanford, California 94305
(415) 497-5801

MILLER, George, Ph.D. [Term expiring]
The Rockefeller University
1230 York Avenue
New York, New York 10021
(212) 360-1801

MOHLER, William C., M.D.
 Associate Director
 Division of Computer Research and Technology
 National Institutes of Health
 Building 12A, Room 3033
 9000 Rockville Pike
 Bethesda, Maryland 20014
 (301) 496-1168

MYERS, Jack D., M.D.
 School of Medicine
 Scaife Hall, 1291
 University of Pittsburgh
 Pittsburgh, Pennsylvania 15261
 (412) 624-2649

REDDY, D.R., Ph.D. [Term expiring]
 Department of Computer Science
 Carnegie-Mellon University
 Pittsburgh, Pennsylvania
 (412) 621-2600, Ext. 149

SAFIR, Aran, M.D.
 Department of Ophthalmology
 Mount Sinai School of Medicine
 City University of New York
 Fifth Avenue and 100th Street
 New York, New York 10029
 (212) 369-4721

STANFORD COMMUNITY ADVISORY COMMITTEE:

LEDERBERG, Dr. Joshua (Chairman)

Principal Investigator - SUMEX
Department of Genetics, S331
Stanford University Medical Center
Stanford, California 94305
(415) 497-5801

COHEN, Stanley N., M.D.

Department of Clinical Pharmacology, S169
Stanford University Medical Center
Stanford, California 94305
(415) 497-5315

DJERASSI, Dr. Carl

Department of Chemistry, Stauffer I-106
Stanford University
Stanford, California 94305
(415) 497-2783

FEIGENBAUM, Dr. Edward

Serra House
Department of Computer Science
Stanford University
Stanford, California 94305
(415) 497-4878

LEVINTHAL, Dr. Elliott C.

Department of Genetics, S047
Stanford University Medical Center
Stanford, California 94305
(415) 497-5813

Appendix VI

USER INFORMATION - GENERAL BROCHURE

Revised May 1976

Appendix VIIGUIDELINES FOR PROSPECTIVE USERSSUMEX-AIM RESOURCE
INFORMATION FOR POTENTIAL USERS

National users may gain access to the facility resources through an advisory panel for a national program in Artificial Intelligence in Medicine (AIM). The AIM Advisory Group consists of members-at-large of the AI and medical communities, facility users and the Principal Investigator of SUMEX as an ex-officio member. A representative of the National Institutes of Health-Biotechnology Resources Program (NIH-BRP) serves as Executive Secretary.

Under its enabling 5-year grant, the SUMEX-AIM computing resource is allocated to qualified users without fee. This, of course, entails a careful review of the merits and priorities of proposed applications. At the direction of the Advisory Group, expenses related to communications and transportation to allow specific users to visit the facility also may be covered.

USER QUALIFICATIONS

The SUMEX-AIM facility is a community effort, not merely a machine service. Applications for membership are judged on the basis of the following criteria:

- 1) The scientific interest and merit of the proposed research and its relevance to the health research missions of the NIH.
- 2) The congruence of research needs and goals to the AI functions of SUMEX-AIM as opposed to other computing alternatives.
- 3) The user's prospective contributions and role in the community, with respect to computer science, e.g., developing and sharing new systems or applications programs, sharing use of special hardware, etc.
- 4) The user's potential for substantive scientific cooperation with the community, e.g., to share expert knowledge in relevant scientific specialties.
- 5) The quantitative demands for specific elements of the SUMEX-AIM resource, taking account of both mean and ceiling requirements.

In many respects, this requires a different kind of information for judgment of proposals than that required for routine grant applications seeking monetary funding support. Information furnished by users also is indispensable to the SUMEX staff in conducting their planning, reporting and operational functions.

The following questionnaire encompasses the main issues concerning the Advisory Group. However, this should neither obstruct clear and imaginative presentation nor restrict format of the application. The potential user should prepare a statement in his own words using previously published material or other documents where applicable. In this respect, the questionnaire may be most useful as a checklist and reference for finding in other documentation the most cogent replies to the questions raised.

For users mounting complex and especially non-standard systems, the decision to affiliate with SUMEX may entail a heavy investment that would be at risk if the arrangement were suddenly terminated. The Advisory Group endeavors to follow a responsible and sensitive policy along these lines--one reason for cautious deliberation; and even in the harshest contingencies, it will make every effort to facilitate graceful entry and departure of qualified users. Conversely, it must have credible information about thoughtful plans for long-term requirements including eventual alternatives to SUMEX-AIM. SUMEX-AIM is a research resource, not an operational vehicle for health care. Many programs are expected to be investigated, developed and demonstrated on SUMEX-AIM with spinoffs for practical implementation on other systems. In some cases, the size, scope and probable validation of clinical trials would preclude their being undertaken on SUMEX-AIM as now constituted. Please be as explicit as possible in your plans for such outcomes.

Applicants, therefore, should submit:

- 1) One to two-page outline of the proposal.
- 2) Response to questionnaire, cross-referenced to supporting documents where applicable.
- 3) Supporting documents.
- 4) List of submitted materials, cross-referenced.

We would welcome a draft (2 copies) of your submission for informal comment if you so desire. However, for formal consideration by the SUMEX-AIM Advisory Group, please submit 13 copies of the material requested above in final form.

Elliott Levinthal, Ph.D.
AIM User Liaison
SUMEX-AIM Computer Project
c/o Department of Genetics, S047
Stanford University Medical Center
Stanford, California 94305
Telephone: (415) 497-5813

May, 1976

SUMEX-AIM RESOURCE

QUESTIONNAIRE FOR POTENTIAL USERS

Please provide either a brief reply to the following or cite supporting documents.

A) MEDICAL AND COMPUTER SCIENCE GOALS

- 1) Describe the proposed research to be undertaken on the SUMEX-AIM resource.
- 2) How is this research presently supported? Please identify application and award statements in which the contingency of SUMEX-AIM availability is indicated. What is the current status of any application for grant support of related research by any federal agency? Please note if you have received notification of any disapproval or approval, pending funding, within the past three years. Budgetary information should be furnished where it concerns operating costs and personnel for computing support. Please furnish any contextual information concerning previous evaluation of your research plans by other scientific review groups.
- 3) What is the relevance of your research to the AI approach of SUMEX-AIM as opposed to other computing alternatives?

B) COLLABORATIVE COMMUNITY BUILDING

- 1) Will the programs designed in your research efforts have some possible general application to problems analogous to that research?
- 2) What application programs already publically available can you use in your research? Are these available on SUMEX-AIM or elsewhere?
- 3) What opportunities or difficulties do you anticipate with regard to making available your programs to other collaborators within a reasonable interval of publication of your work?
- 4) Are you interested in discussing with the SUMEX staff possible ways in which other artificial-intelligence research capabilities might interrelate with your work?
- 5) If approved as a user, would you advise us regarding collaborative opportunities similar to yours with other investigators in your field?

C) HARDWARE AND SOFTWARE REQUIREMENTS

- 1) What computer facilities are you now using in connection with your research or do you have available at your institution? In what respect do these not meet your research requirements?

- 2) What languages do you either use or wish to use? Will your research require the addition of major system programs or languages to the system? Will you maintain them? If you are committed to systems not now maintained at SUMEX, what effort would be required for conversion to and maintenance on the PDP-10 - TENEX system? What are the merits of the alternative plan of converting your application programs to one of the already available standards? Would the latter facilitate the objectives of Part B), Collaborative Community Building?
- 3) Can you estimate your requirements for CPU utilization and disk space? What time of day will your CPU utilization occur? Would it be convenient or possible for you to use the system during off-peak periods? Please indicate (as best you can) the basis for these estimates and the consequences of various levels of restriction or relaxation of access to different resources. SUMEX-AIM's tangible resources can be measured in terms of:
 - a) CPU cycles.
 - b) Connect time and communications.
 - c) User terminals (In special cases these may be supported by SUMEX-AIM.).
 - d) Disk space.
 - e) Off-line media-printer outputs, tapes (At most, limited quantities to be mailed.).

Can you estimate your requirements? With respect to a) and b), there are loading problems during the daily cycle.--Can you indicate the relative utility of prime-time (0900-1600 PST) vs. off-peak access?

- 4) What are your communication plans (TYMNET, ARPANET, other)? How will your communication and terminal costs be met? See following note concerning network connections to SUMEX-AIM.
- 5) If this is a development project, please indicate your long-term plans for software implementation in an applied context keeping in mind the research mission of SUMEX-AIM.

Our procedures are still evolving, and we welcome your suggestions about this framework for exchanging information. Needless to say, each question should be qualified a) "insofar as relevant to your proposal", and b) "to the extent of available information".

Please do not force a reply to a question that seems inappropriate. We prefer that you label it as such so that it can be dealt with properly in future dialogue.

Above all, we are eager to work with potential users in any way that would help minimize bureaucratic burdens and still permit a responsible regard for our accountability both to the NIH and the public. Please do not hesitate to address the substance of these requirements in the format most applicable to you.

NETWORK CONNECTIONS TO SUMEX-AIM

TYMNET

Attached is a list of available TYMNET nodes and associated telephone numbers. The cost to users of using TYMNET is the telephone charge from user location to the nearest TYMNET node. This is available only for communication to SUMEX-AIM and not for other facilities that may be connected to TYMNET. In some cases, there are "foreign exchanges" set up by users. These may offer less expensive communication. Details of these possibilities can best be learned by calling the nearest TYMNET node. The telephone company can provide information on comparative costs of leased lines, toll charges, etc. The initial capital investment for TYMNET installation as well as login and hourly charges is provided by SUMEX-AIM. Standard usage charges on TYMNET are approximately \$3/connect-hour.

ARPANET

SUMEX-AIM is connected to the ARPANET. Our name is SUMEX-AIM; our nickname is AIM. We support the new TELNET protocol. Our network address is decimal 56, octal 70. This provides convenient access for ARPANET Hosts and Associates and those who have accounts with ARPANET.